ESnet4: Networking for the Future of DOE Science

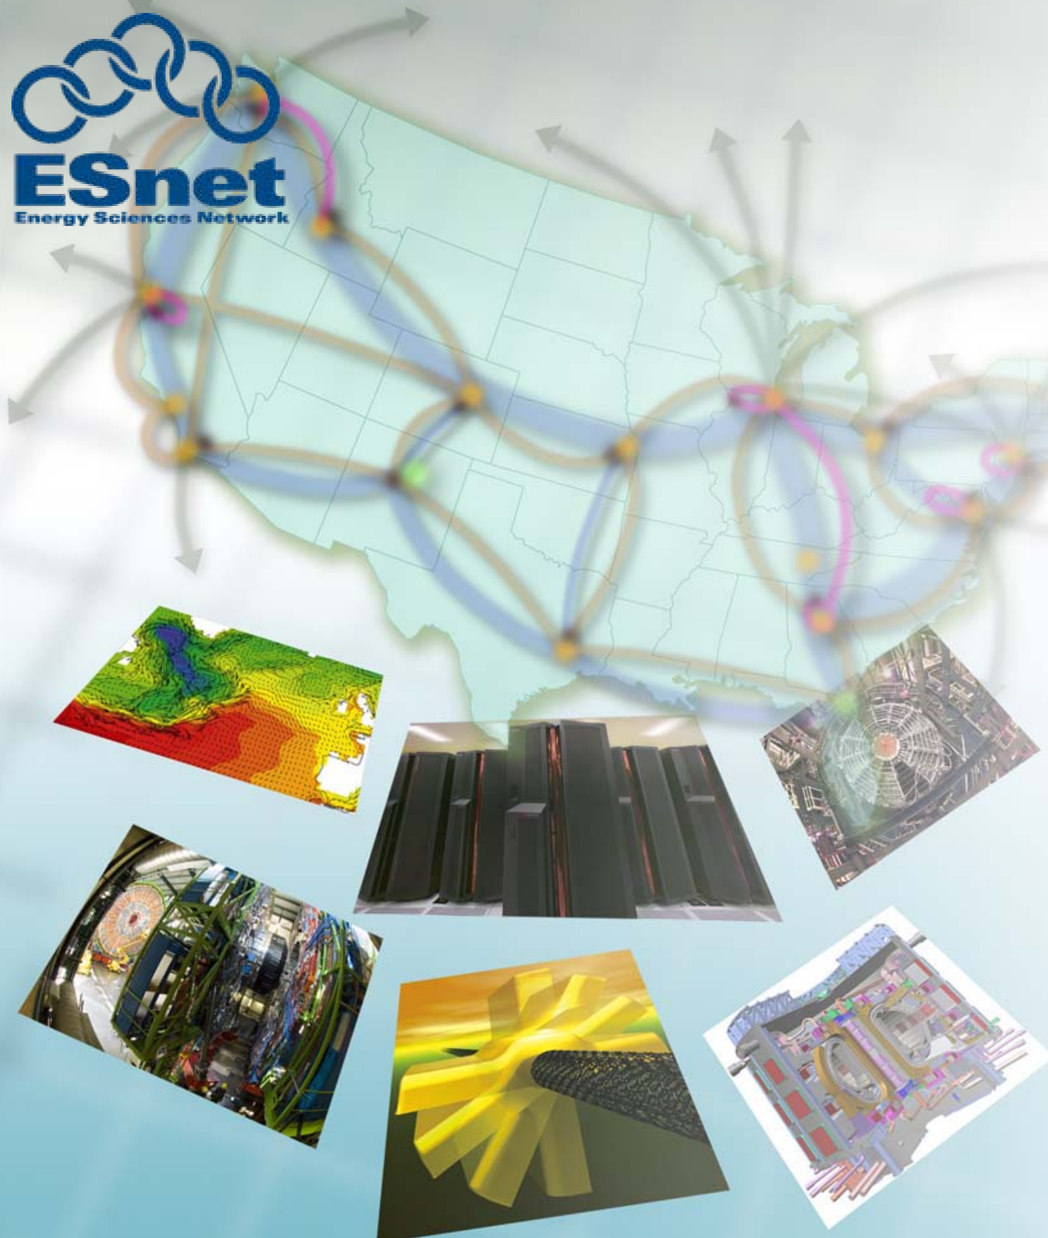

International ICFA Workshop on Digital Divide Issues for Global e-Science
October 25, 2007

Joe Metzger
Senior Engineer

Energy Sciences Network
Lawrence Berkeley National Laboratory

metzger@es.net, www.es.net

DOE's Office of Science: Enabling Large-Scale Science

- The Office of Science (SC) is the single ***largest supporter of basic research in the physical sciences in the United States***, … providing more than 40 percent of total funding … for the Nation's research programs in high-energy physics, nuclear physics, and fusion energy sciences. (http://www.science.doe.gov) – SC funds 25,000 PhDs and PostDocs
- A primary mission of SC's National Labs is to ***build and operate very large scientific instruments*** - particle accelerators, synchrotron light sources, very large supercomputers - that generate massive amounts of data and involve very large, distributed collaborations
- **ESnet - the Energy Sciences Network - is an SC program whose *primary mission is to enable the large-scale science* of the Office of Science that depends on**:
  - Sharing of massive amounts of data
  - Supporting thousands of collaborators world-wide
  - Distributed data processing
  - Distributed data management
  - Distributed simulation, visualization, and computational steering
  - Collaboration with the US and International Research and Education community
- In addition to the National Labs, ESnet serves much of the rest of DOE, including NNSA – about 75,000-100,000 users in total

Office of Science US Community

Drives ESnet Design for Domestic Connectivity

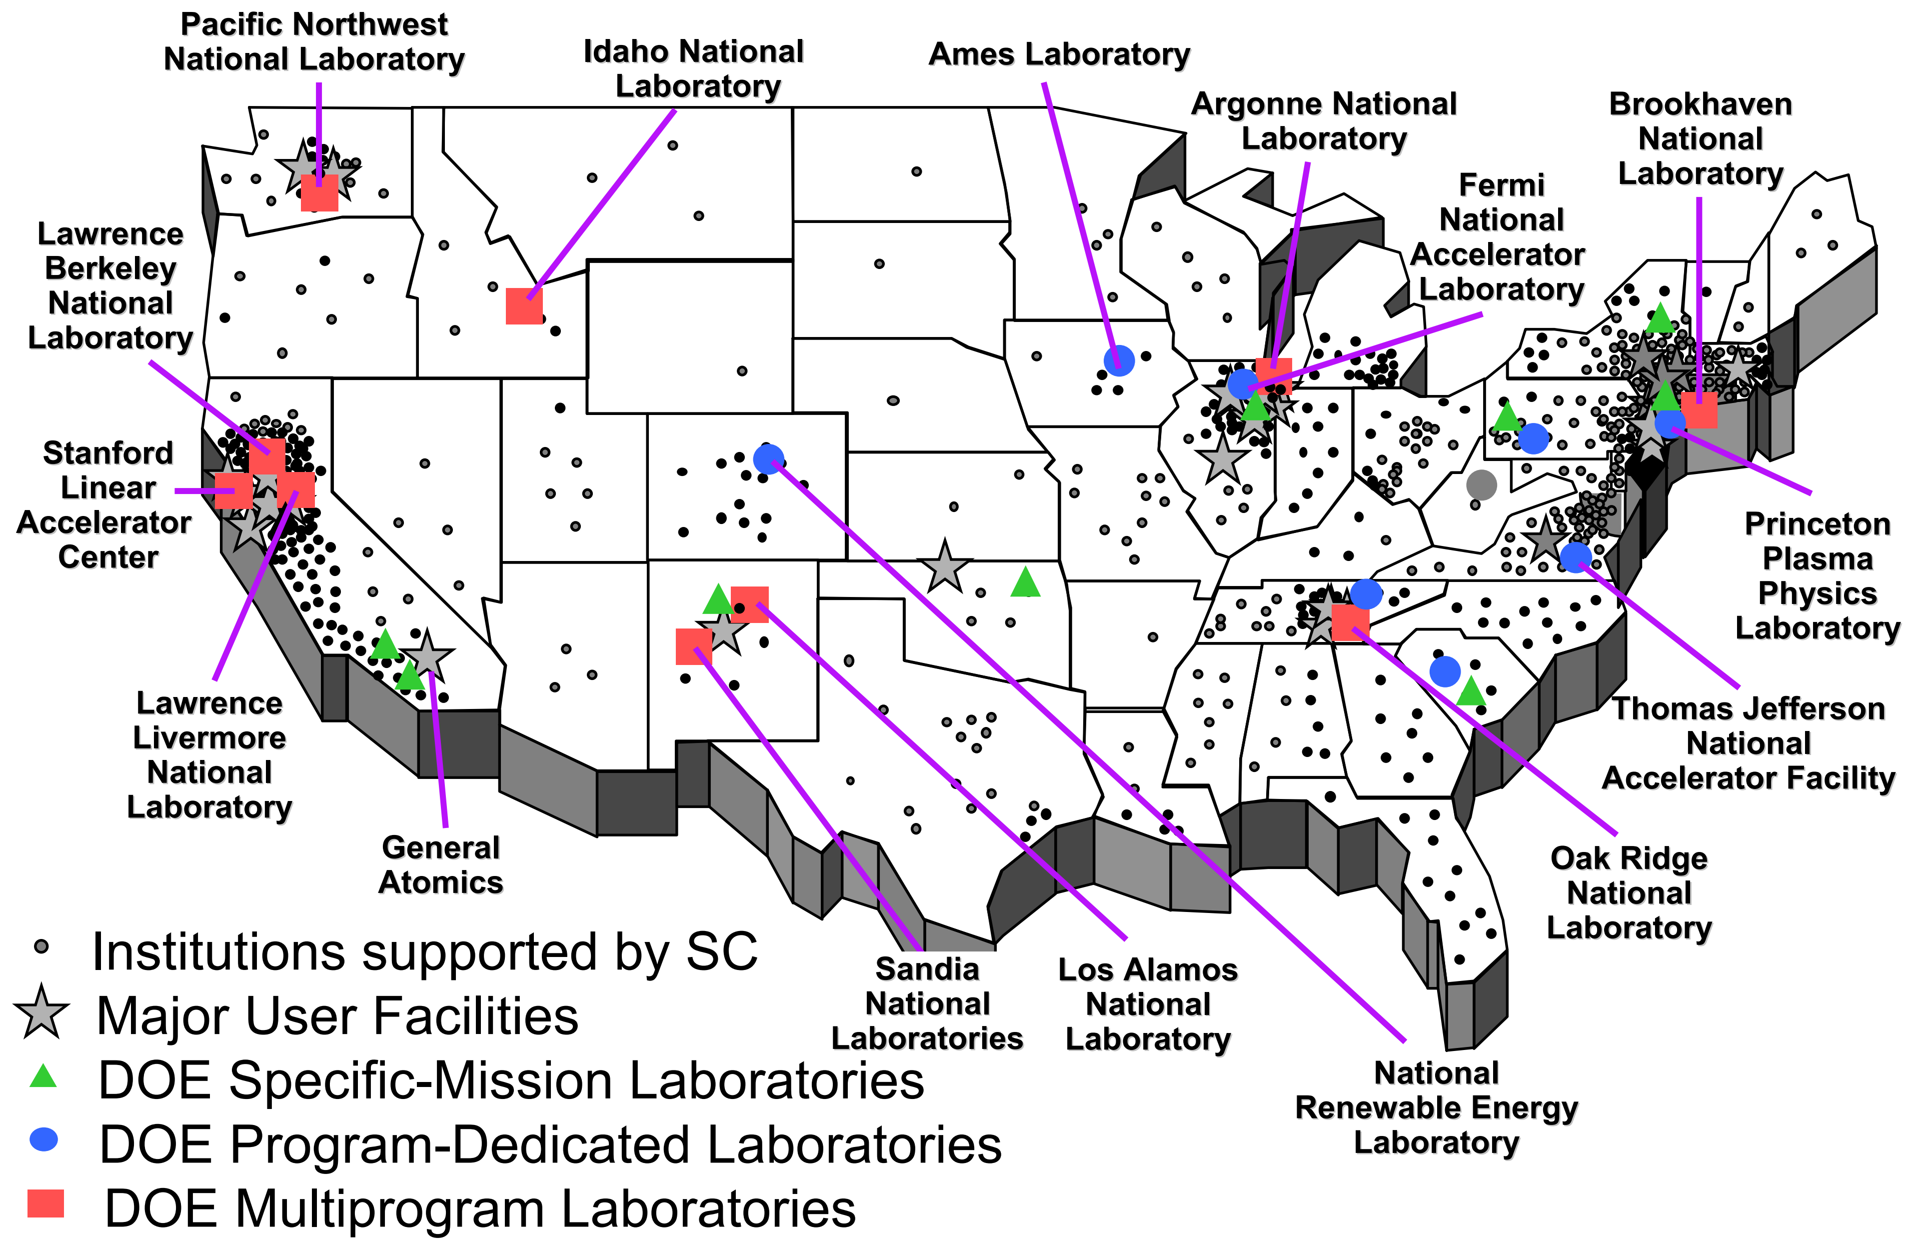

ESnet as a Large-Scale Science Data Transport Network

<u>A few hundred of the 25,000,000 hosts that the Labs connect to every month account for 50% of all ESnet traffic</u> (which is all science data - primarily high energy physics, nuclear physics, and climate at this point)

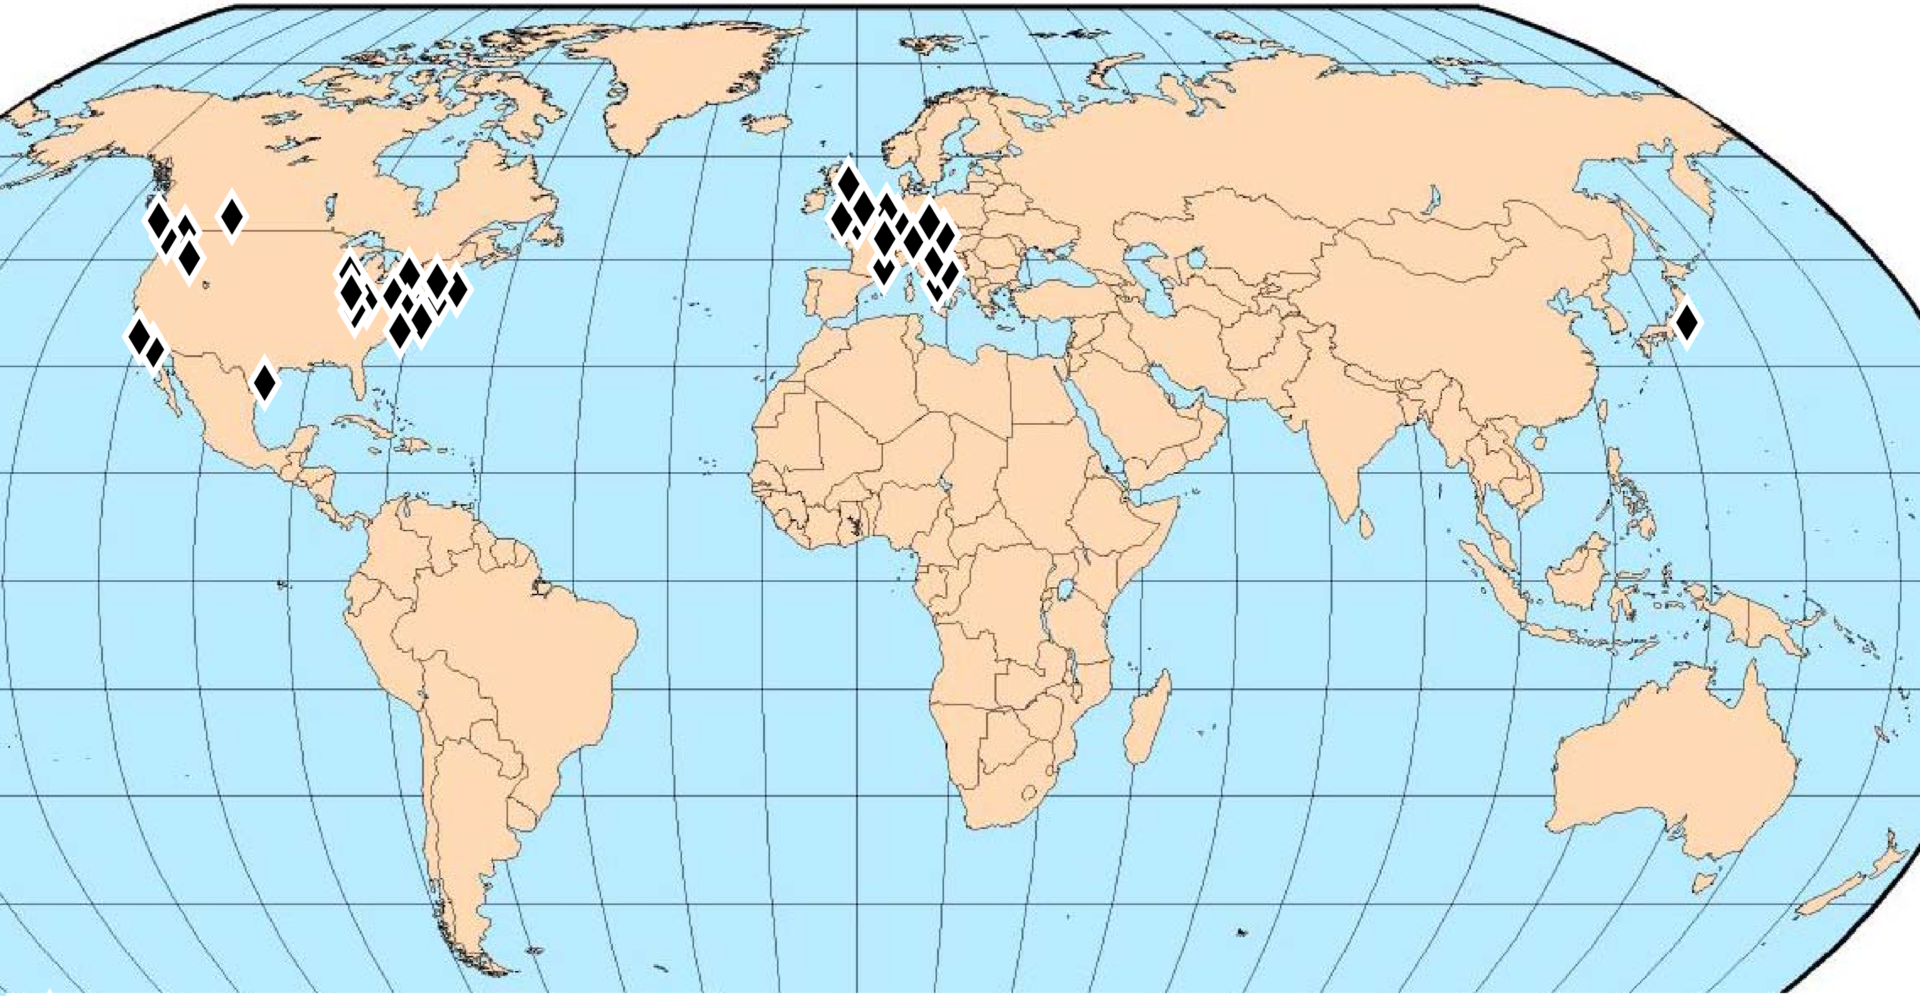

♦= the R&E source or destination of ESnet's top 100 sites (all R&E) (the DOE Lab destination or source of each flow is not shown)

What ESnet Is

- A large-scale IP network built on a national circuit infrastructure with high-speed connections to all major US and international research and education (R&E) networks
- An organization of 30 professionals structured for the service
- An operating entity with an FY06 budget of $26.6M
- A tier 1 ISP (direct peerings will all major networks)
- The primary DOE network provider
  - Provides production Internet service to all of the major DOE Labs* and most other DOE sites
  - Based on DOE Lab populations, it is estimated that between 50,000 - 100,000 users depend on ESnet for global Internet access
    - additionally, each year more than 18,000 non-DOE researchers from universities, other government agencies, and private industry use Office of Science facilities

* PNNL supplements its ESnet service with commercial service

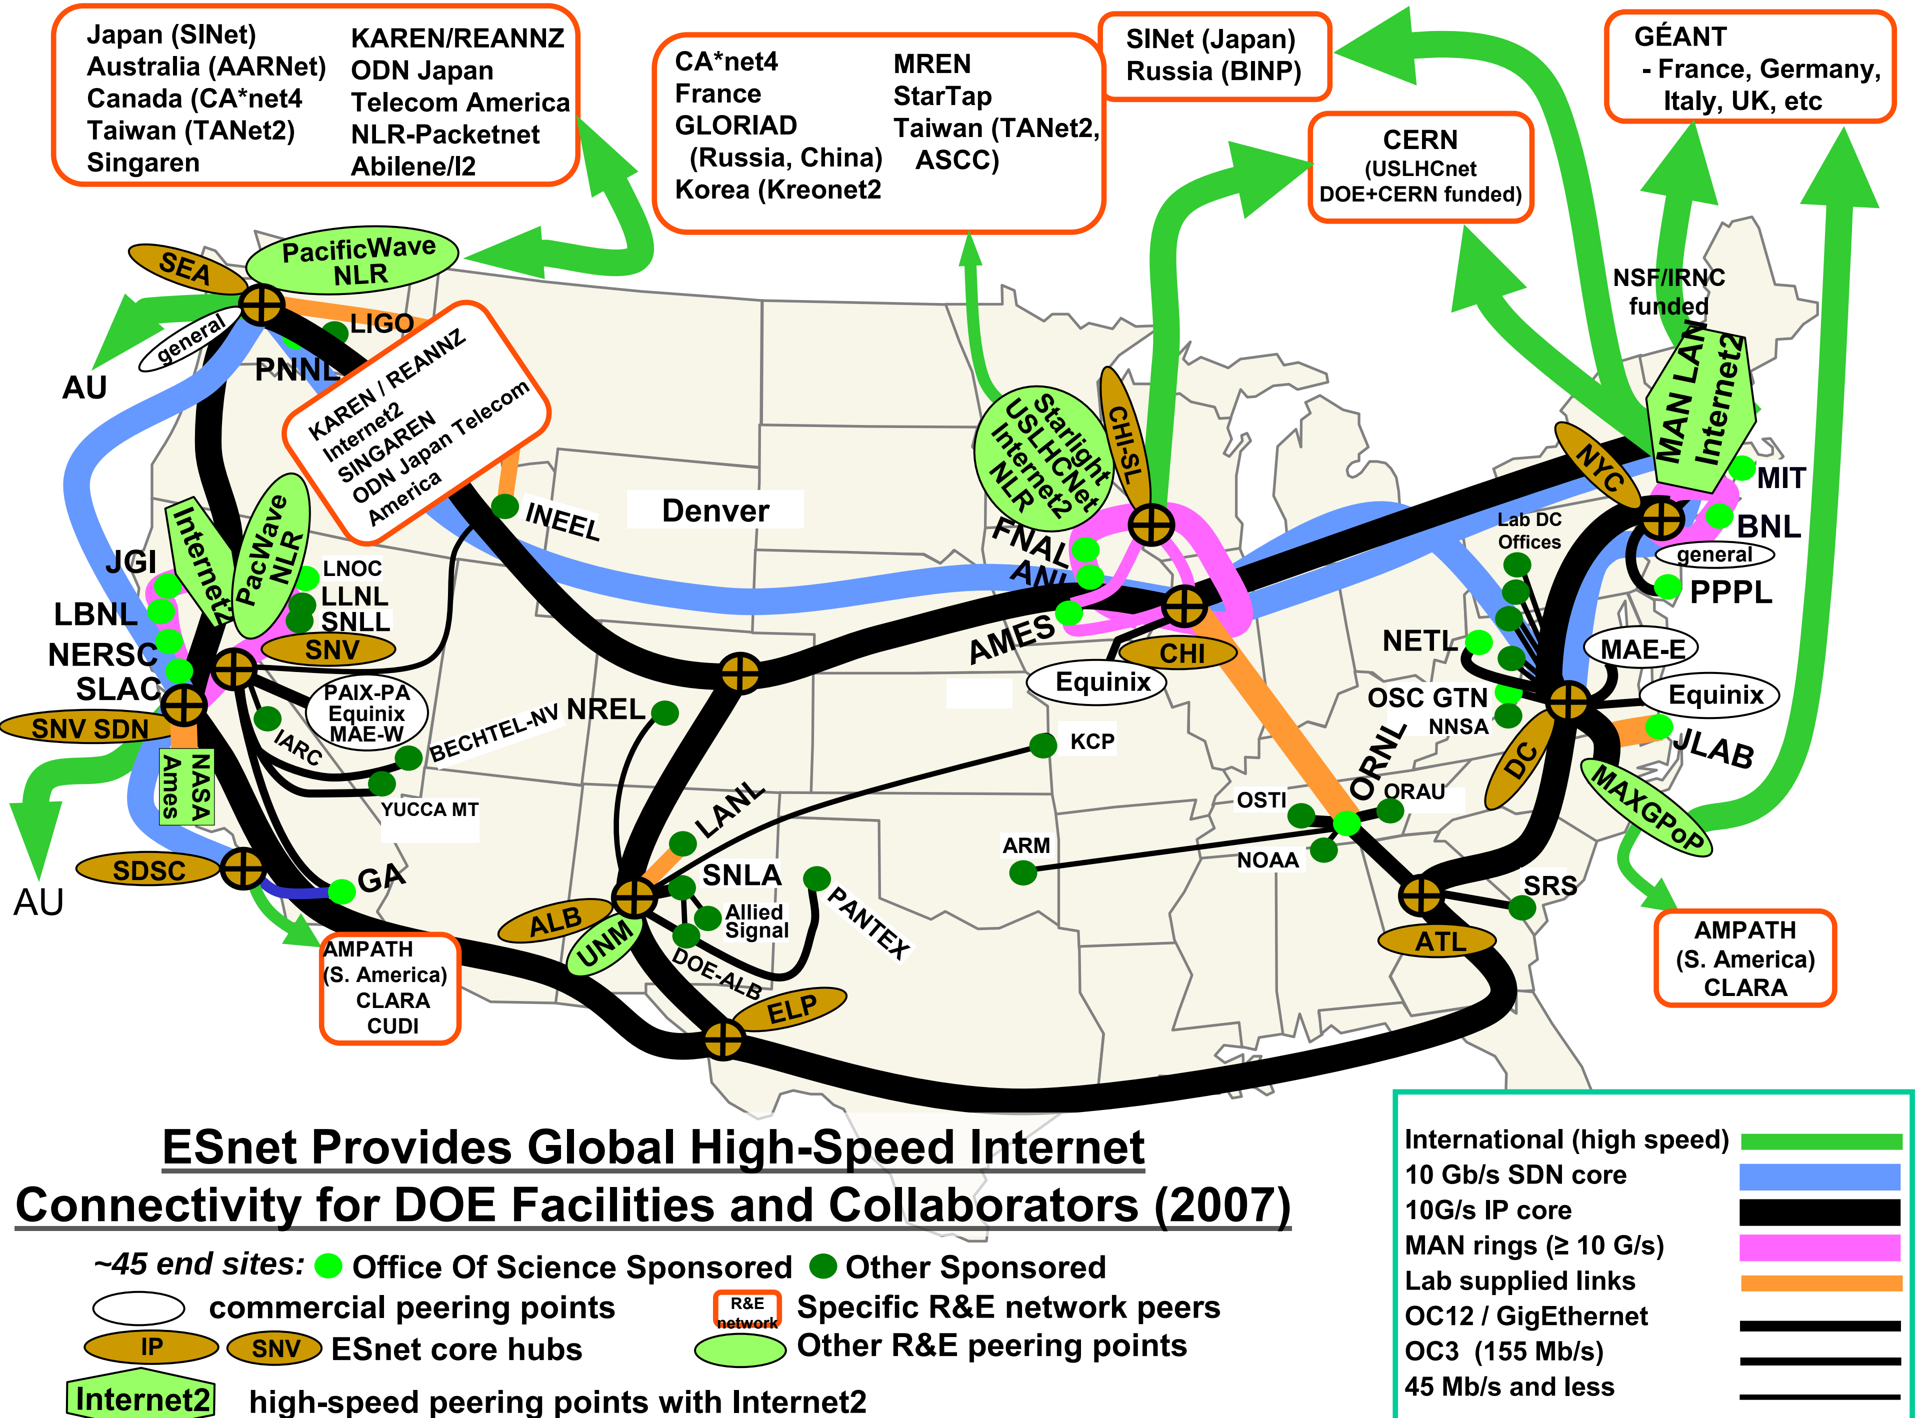

ESnet Provides Global High-Speed Internet Connectivity for DOE Facilities and Collaborators (2007)

~45 end sites: Office Of Science Sponsored — Other Sponsored

commercial peering points — R&E network: Specific R&E network peers

IP, SNV: ESnet core hubs — Other R&E peering points

Internet2: high-speed peering points with Internet2

ESnet's Place in U. S. and International Science

- ESnet, Internet2/Abilene, and National Lambda Rail (NLR) provide most of the nation's transit networking for basic science
  - Abilene provides national transit networking for most of the US universities by interconnecting the regional networks (mostly via the GigaPoPs)
  - ESnet provides national transit networking and ISP service for the DOE Labs
  - NLR provides various science-specific and network R&D circuits
- GÉANT plays a role in Europe similar to Abilene and ESnet in the US – it interconnects the European National Research and Education Networks (NRENs), to which the European R&E sites connect
  - A GÉANT operated, NSF funded link currently carries all non-LHC ESnet traffic to Europe, and this is a significant fraction of all ESnet traffic

Operating *Science Mission* Critical Infrastructure

- ESnet is a visible and critical piece of DOE science infrastructure
  - if ESnet fails,10s of thousands of DOE and University users know it within minutes if not seconds
- Requires high reliability and high operational security in the systems that are integral to the operation and management of the network
  - Secure and redundant mail and Web systems are central to the operation and security of ESnet
    - trouble tickets are by email
    - engineering communication by email
    - engineering database interfaces are via Web
  - Secure network access to Hub routers
  - Backup secure telephone modem access to Hub equipment
  - 24x7 help desk and 24x7 on-call network engineer

    trouble@es.net (end-to-end problem resolution)

➢A Changing Science Environment is the Key Driver of the Next Generation ESnet

- Large-scale collaborative science – big facilities, massive data, thousands of collaborators – is now a significant aspect of the Office of Science (“SC”) program
- SC science community is almost equally split between Labs and universities
  - SC facilities have users worldwide
- Very large international (non-US) facilities (e.g. LHC and ITER) and international collaborators are now a key element of SC science
- Distributed systems for data analysis, simulations, instrument operation, etc., are essential and are now common (in fact dominate data analysis that now generates 50% of all ESnet traffic)

Planning and Building ESnet4

- Requirements are primary drivers for ESnet – science focused

- Sources of Requirements

1. Office of Science (SC) Program Managers
    - The Program Offices Requirements Workshops
        - 2007
            - » Basic Energy Science
            - » Biological & Environmental Research
        - 2008
            - » Fusion Energy Science
            - » Nuclear Physics
        - 2009
            - » High Energy Physics
            - » Advance Scientific Computing Research
2. Direct gathering through interaction with science users of the network
    - Example case studies (from 2005/2006 update)
        - Magnetic Fusion
        - Large Hadron Collider (LHC)
        - Climate Modeling
        - Spallation Neutron Source
3. Observation of the network

Science Networking Requirements Aggregation Summary

Science Drivers Science Areas / Facilities	End2End Reliability	Connectivity	Today End2End Band width	5 years End2End Band width	Traffic Characteristics	Network Services
Magnetic Fusion Energy	**99.999%** **(Impossible without full redundancy)**	• **DOE sites** • **US Universities** • **Industry**	**200+ Mbps**	**1 Gbps**	• **Bulk data** • **Remote control**	• **Guaranteed bandwidth** • **Guaranteed QoS** • **Deadline scheduling**
NERSC and ACLF	**-**	• **DOE sites** • **US Universities** • **International** • **Other ASCR supercomputers**	**10 Gbps**	**20 to 40 Gbps**	• **Bulk data** • **Remote control** • **Remote file system sharing**	• **Guaranteed bandwidth** • **Guaranteed QoS** • **Deadline Scheduling** • **PKI / Grid**
NLCF	**-**	• **DOE sites** • **US Universities** • **Industry** • **International**	**Backbone Band width parity**	**Backbone band width parity**	• **Bulk data** • **Remote file system sharing**	
Nuclear Physics (RHIC)	**-**	• **DOE sites** • **US Universities** • **International**	**12 Gbps**	**70 Gbps**	• **Bulk data**	• **Guaranteed bandwidth** • **PKI / Grid**
Spallation Neutron Source	**High** **(24x7 operation)**	• **DOE sites**	**640 Mbps**	**2 Gbps**	• **Bulk data**	

Science Network Requirements Aggregation Summary

Science Drivers Science Areas / Facilities	End2End Reliability	Connectivity	Today End2End Band width	5 years End2End Band width	Traffic Characteristics	Network Services
Advanced Light Source	-	• DOE sites • US Universities • Industry	1 TB/day 300 Mbps	5 TB/day 1.5 Gbps	• Bulk data • Remote control	• Guaranteed bandwidth • PKI / Grid
Bioinformatics	-	• DOE sites • US Universities	625 Mbps 12.5 Gbps in two years	250 Gbps	• Bulk data • Remote control • Point-to-multipoint	• Guaranteed bandwidth • High-speed multicast
Chemistry / Combustion	-	• DOE sites • US Universities • Industry	-	10s of Gigabits per second	• Bulk data	• Guaranteed bandwidth • PKI / Grid
Climate Science	-	• DOE sites • US Universities • International	-	5 PB per year 5 Gbps	• Bulk data • Remote control	• Guaranteed bandwidth • PKI / Grid
		Immediate Requirements and Drivers				
High Energy Physics (LHC)	99.95+% (Less than 4 hrs/year)	• US Tier1 (FNAL, BNL) • US Tier2 (Universities) • International (Europe, Canada)	10 Gbps	60 to 80 Gbps (30-40 Gbps per US Tier1)	• Bulk data • Coupled data analysis processes	• Guaranteed bandwidth • Traffic isolation • PKI / Grid

Are These Estimates Realistic? YES!

FNAL outbound CMS traffic for 4 months, to Sept. 1, 2007

Max= 8.9 Gb/s (1064 MBy/s of data), Average = 4.1 Gb/s (493 MBy/s of data)

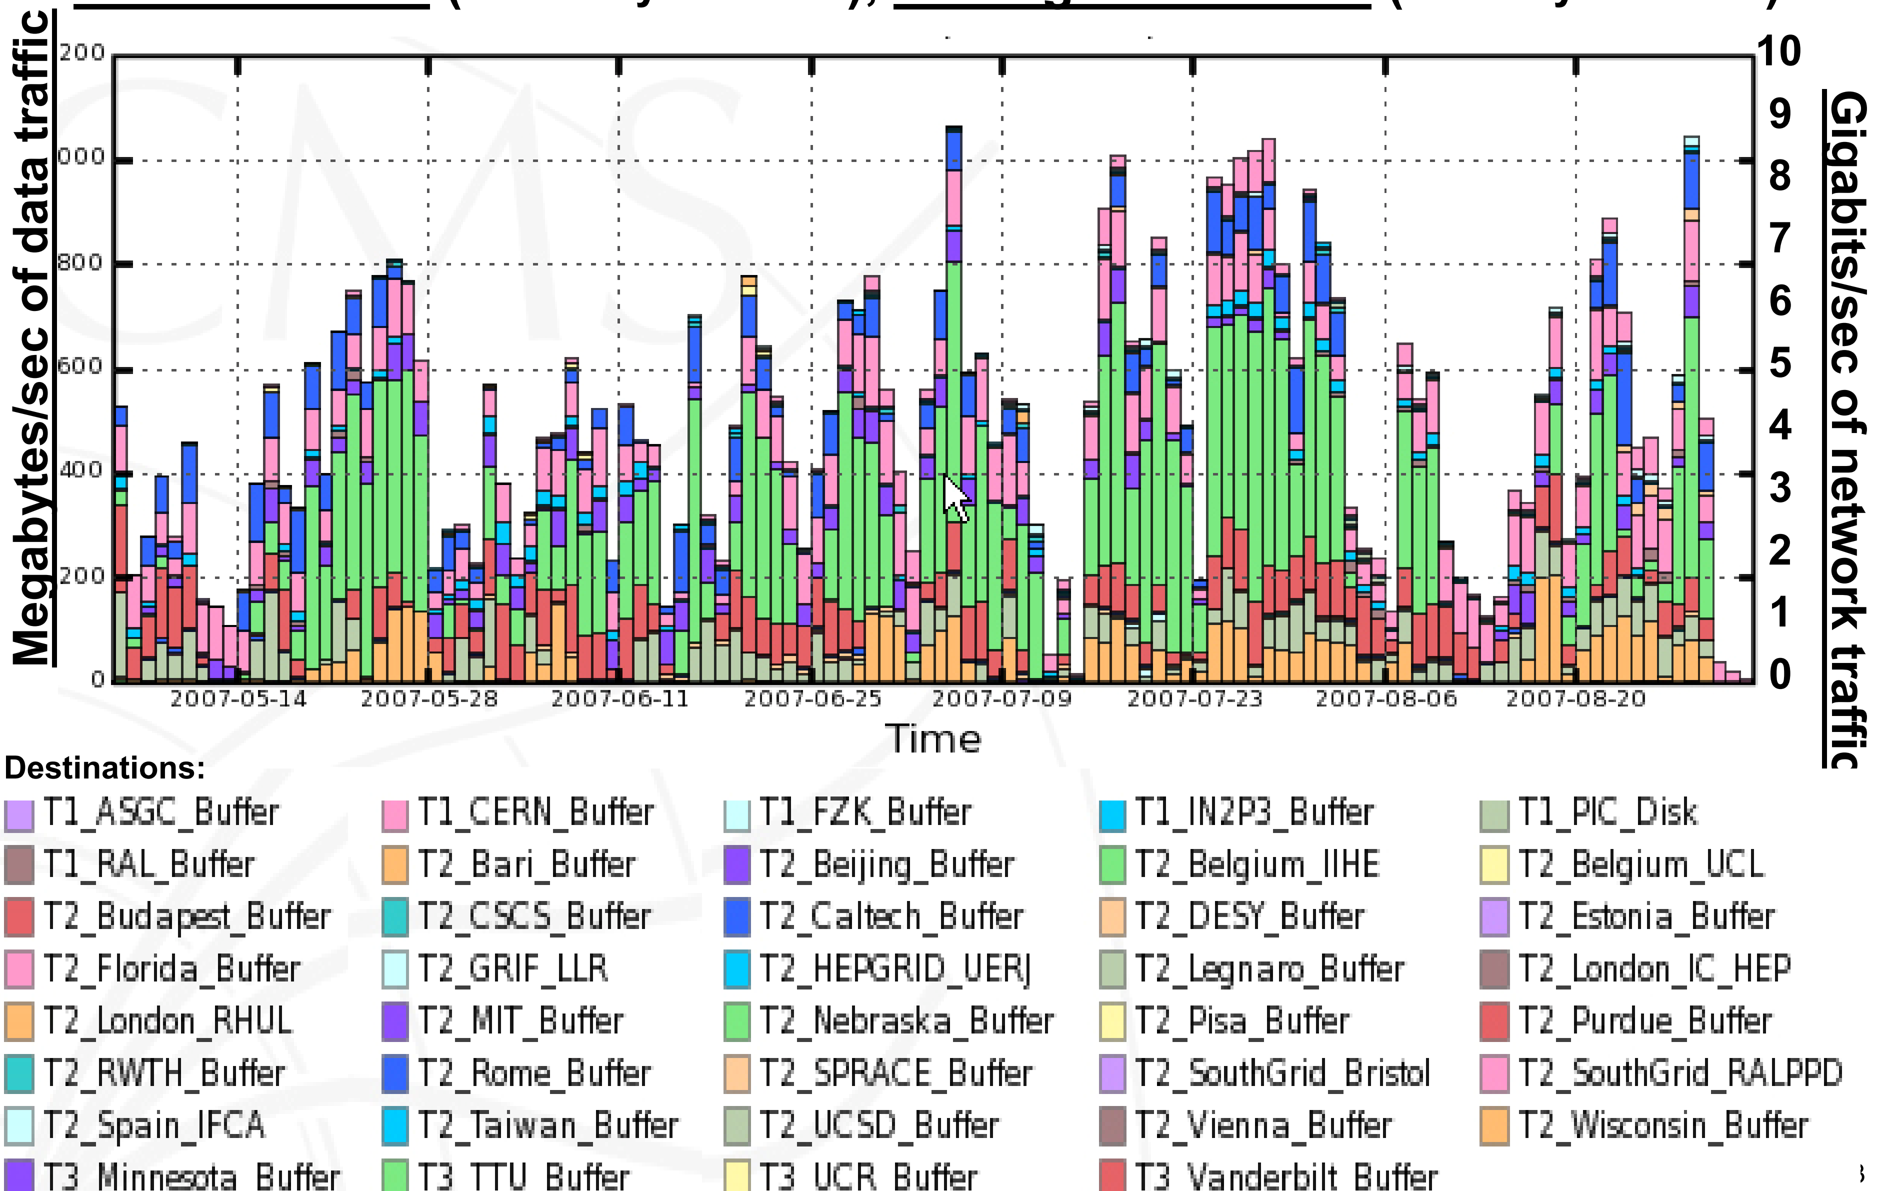

Large-Scale Data Analysis Systems (Typified by the LHC) have Several Characteristics that Result in Requirements for the Network and its Services

- The ***systems are data intensive and high-performance***, typically moving terabytes a day for months at a time
- The ***system are high duty-cycle***, operating most of the day for months at a time in order to meet the requirements for data movement
- The ***systems are widely distributed*** – typically spread over continental or inter-continental distances
- Such ***systems depend on network performance and availability***, but these characteristics cannot be taken for granted, even in well run networks, when the multi-domain network path is considered
- The applications ***must be able to get guarantees from the network*** that there is adequate bandwidth to accomplish the task at hand
- The applications ***must be able to get information from the network*** that allows graceful failure and auto-recovery and adaptation to unexpected network conditions that are short of outright failure
- *These requirements are generally true for systems with widely distributed components to be reliable and consistent in performing the sustained, complex tasks of large-scale science*

This slide drawn from [ICFA SCIC]

Enabling Large-Scale Science

- These requirements are generally true for systems with widely distributed components to be reliable and consistent in performing the sustained, complex tasks of large-scale science
- Networks must provide communication capability that is service-oriented: configurable, schedulable, predictable, reliable, and informative – and the network and its services must be scalable

3. Observed Evolution of Historical ESnet Traffic Patterns

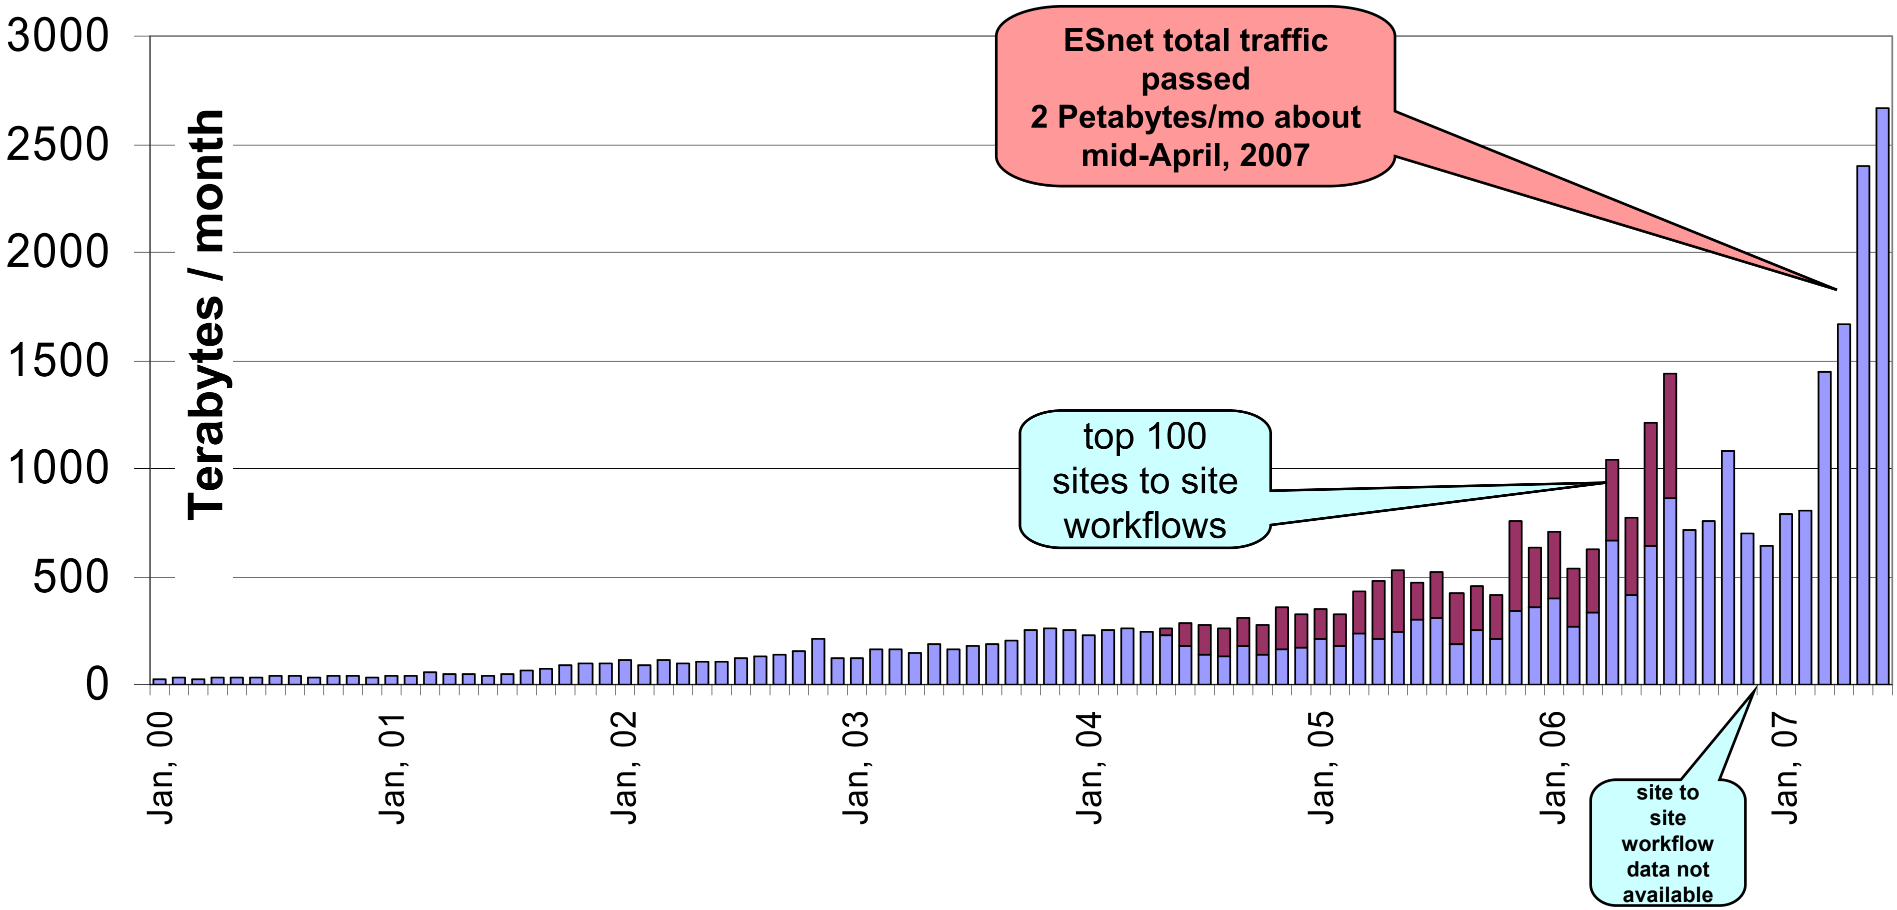

ESnet Monthly Accepted Traffic, January, 2000 – June, 2007

- ESnet is currently transporting more than1 petabyte (1000 terabytes) per month
- More than 50% of the traffic is now generated by the top 100 sites ⇒ large-scale science dominates all ESnet traffic

ESnet Traffic has Increased by **10X Every 47 Months**, on Average, Since 1990

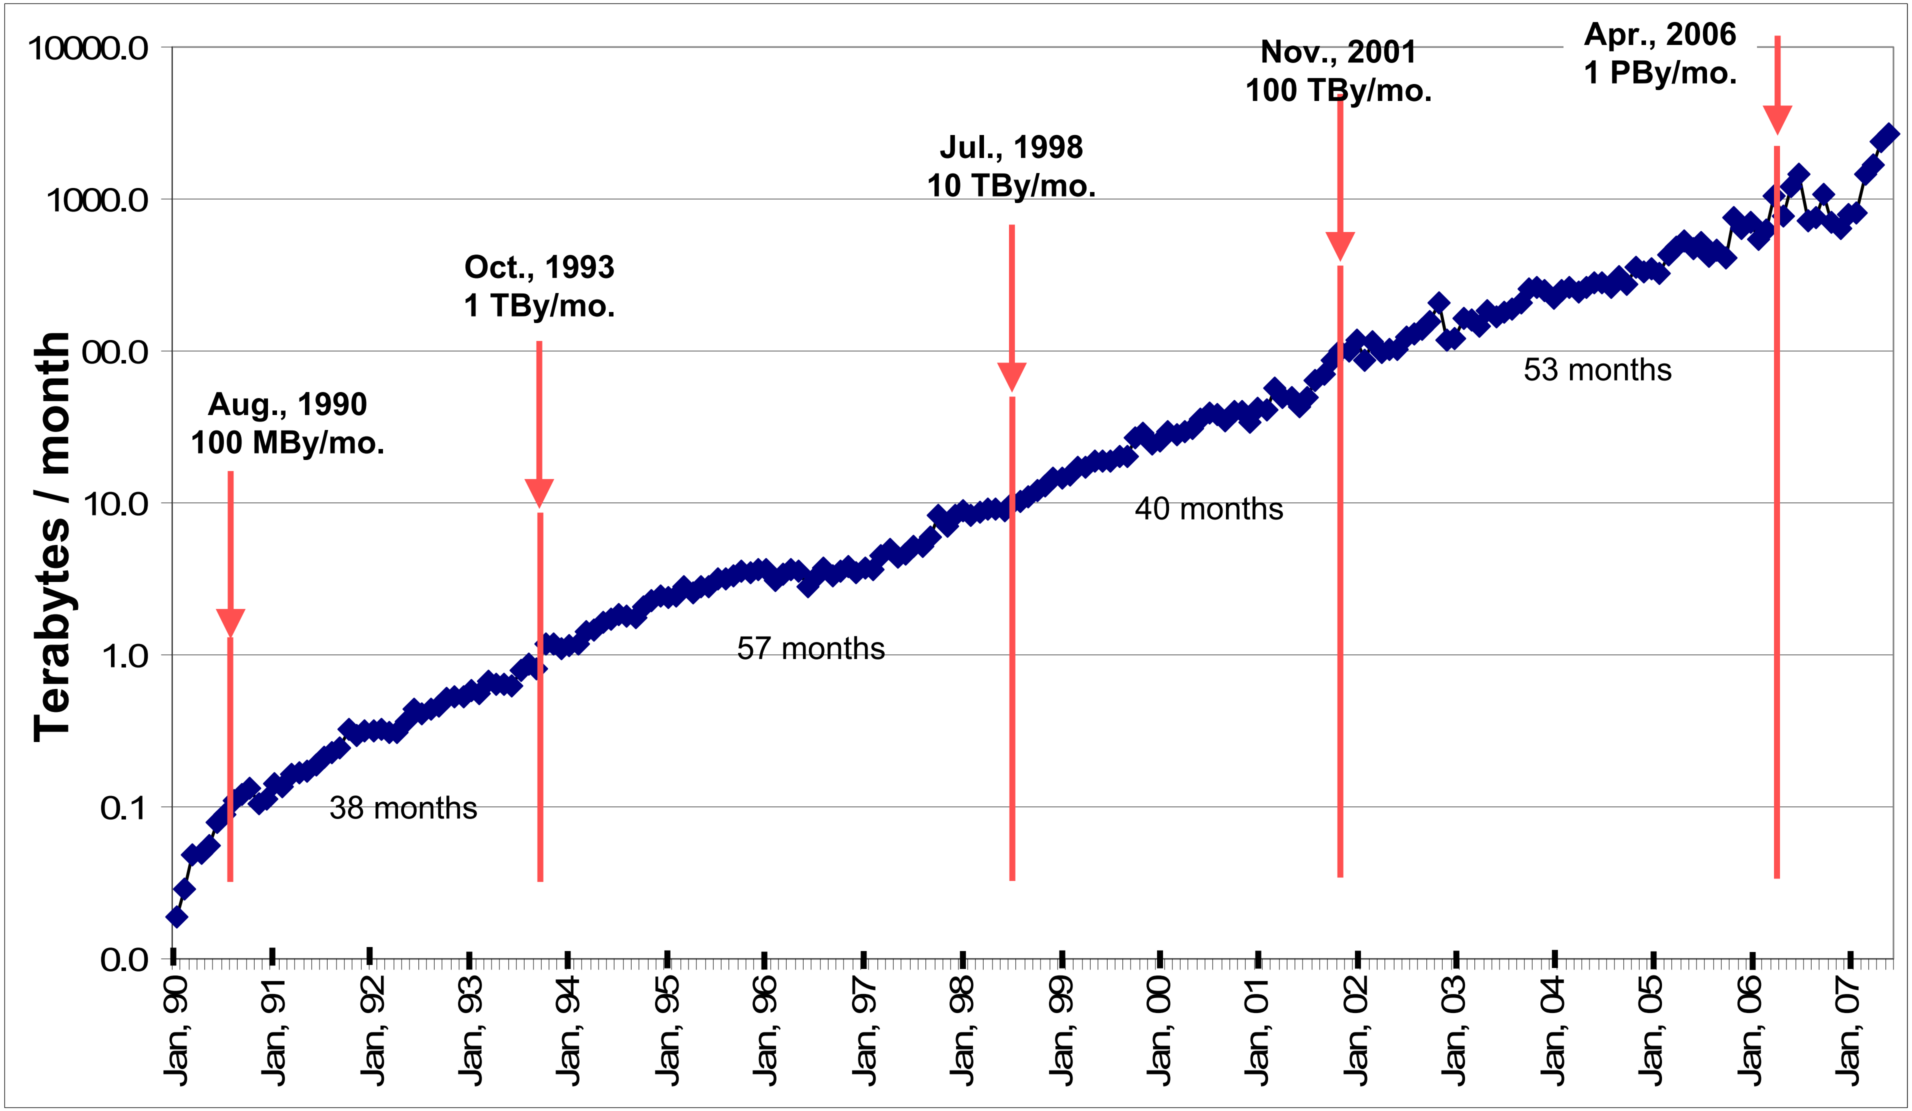

Log Plot of ESnet Monthly Accepted Traffic, January, 1990 – June, 2007

What is the High-Level View of all ESnet Traffic

ESnet Inter-Sector Traffic Summary, May 2007

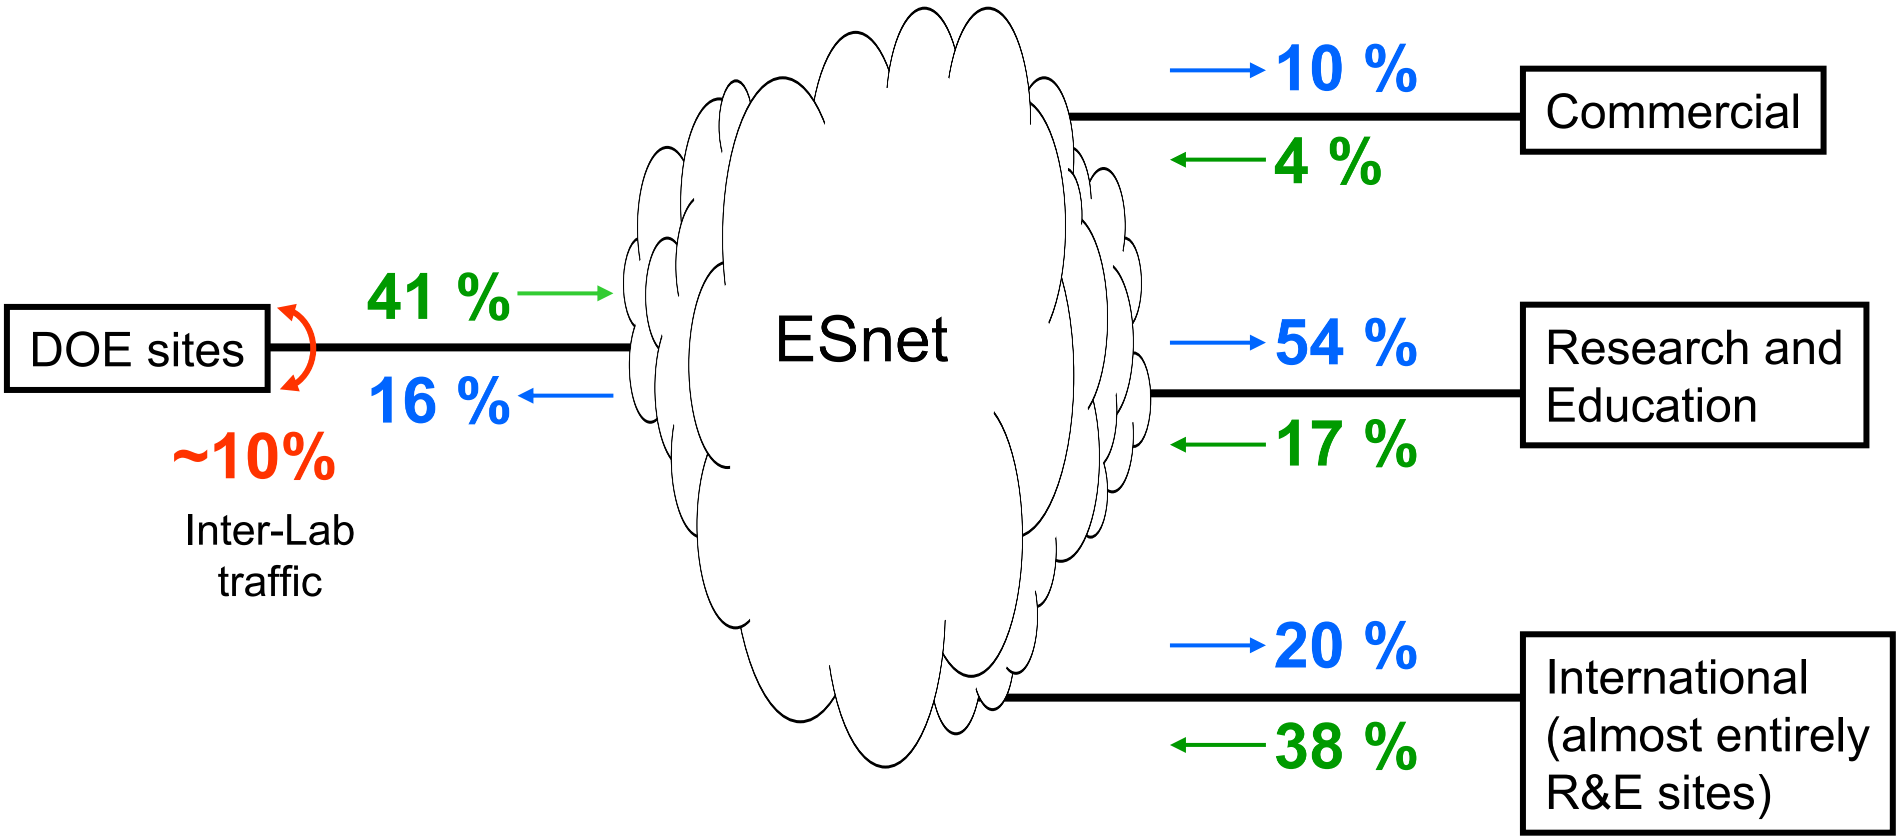

Traffic notes

- *more than 90% of all traffic is Office of Science*

Traffic coming into ESnet = Green
Traffic leaving ESnet = Blue
Traffic between ESnet sites

Requirements from Network Utilization Observation

- In 4 years, we can expect a 10x increase in traffic over current levels *without the addition of production LHC traffic*
  - Nominal average load on busiest backbone links is ~1.5 Gbps today
  - In 4 years that figure will be ~15 Gbps based on current trends
- Measurements of this type are science-agnostic
  - It doesn’t matter who the users are, the traffic load is increasing exponentially
  - Predictions based on this sort of forward projection tend to be conservative estimates of future requirements because they cannot predict new uses

Requirements from Traffic Flow Observations

- Most of ESnet science traffic has a source or sink outside of ESnet
  - Drives requirement for high-bandwidth peering
  - Reliability and bandwidth requirements demand that peering be redundant
  - Multiple 10 Gbps peerings today, must be able to add more bandwidth flexibly and cost-effectively
  - Bandwidth and service guarantees must traverse R&E peerings
    - Collaboration with other R&E networks on a common framework is critical
    - Seamless fabric
- Large-scale science is now the dominant user of the network
  - Satisfying the demands of large-scale science traffic into the future will require a purpose-built, scalable architecture
  - Traffic patterns are different than commodity Internet

Summary of All Requirements To-Date

Requirements from SC Programs:

1A) Provide "consulting" on system / application network tuning

Requirements from science case studies:

2A) Build the ESnet core up to 100 Gb/s within 5 years

2B) Deploy network to accommodate LHC collaborator footprint

2C) Implement network to provide for LHC data path loadings

2D) Provide the network as a service-oriented capability

Requirements from observing traffic growth and change trends in the network:

3A) Provide 15 Gb/s core within four years and 150 Gb/s core within eight years

3B) Provide a rich diversity and high bandwidth for R&E peerings

3C) Economically accommodate a very large volume of circuit-like traffic

Estimated Aggregate Link Loadings, 2007-08

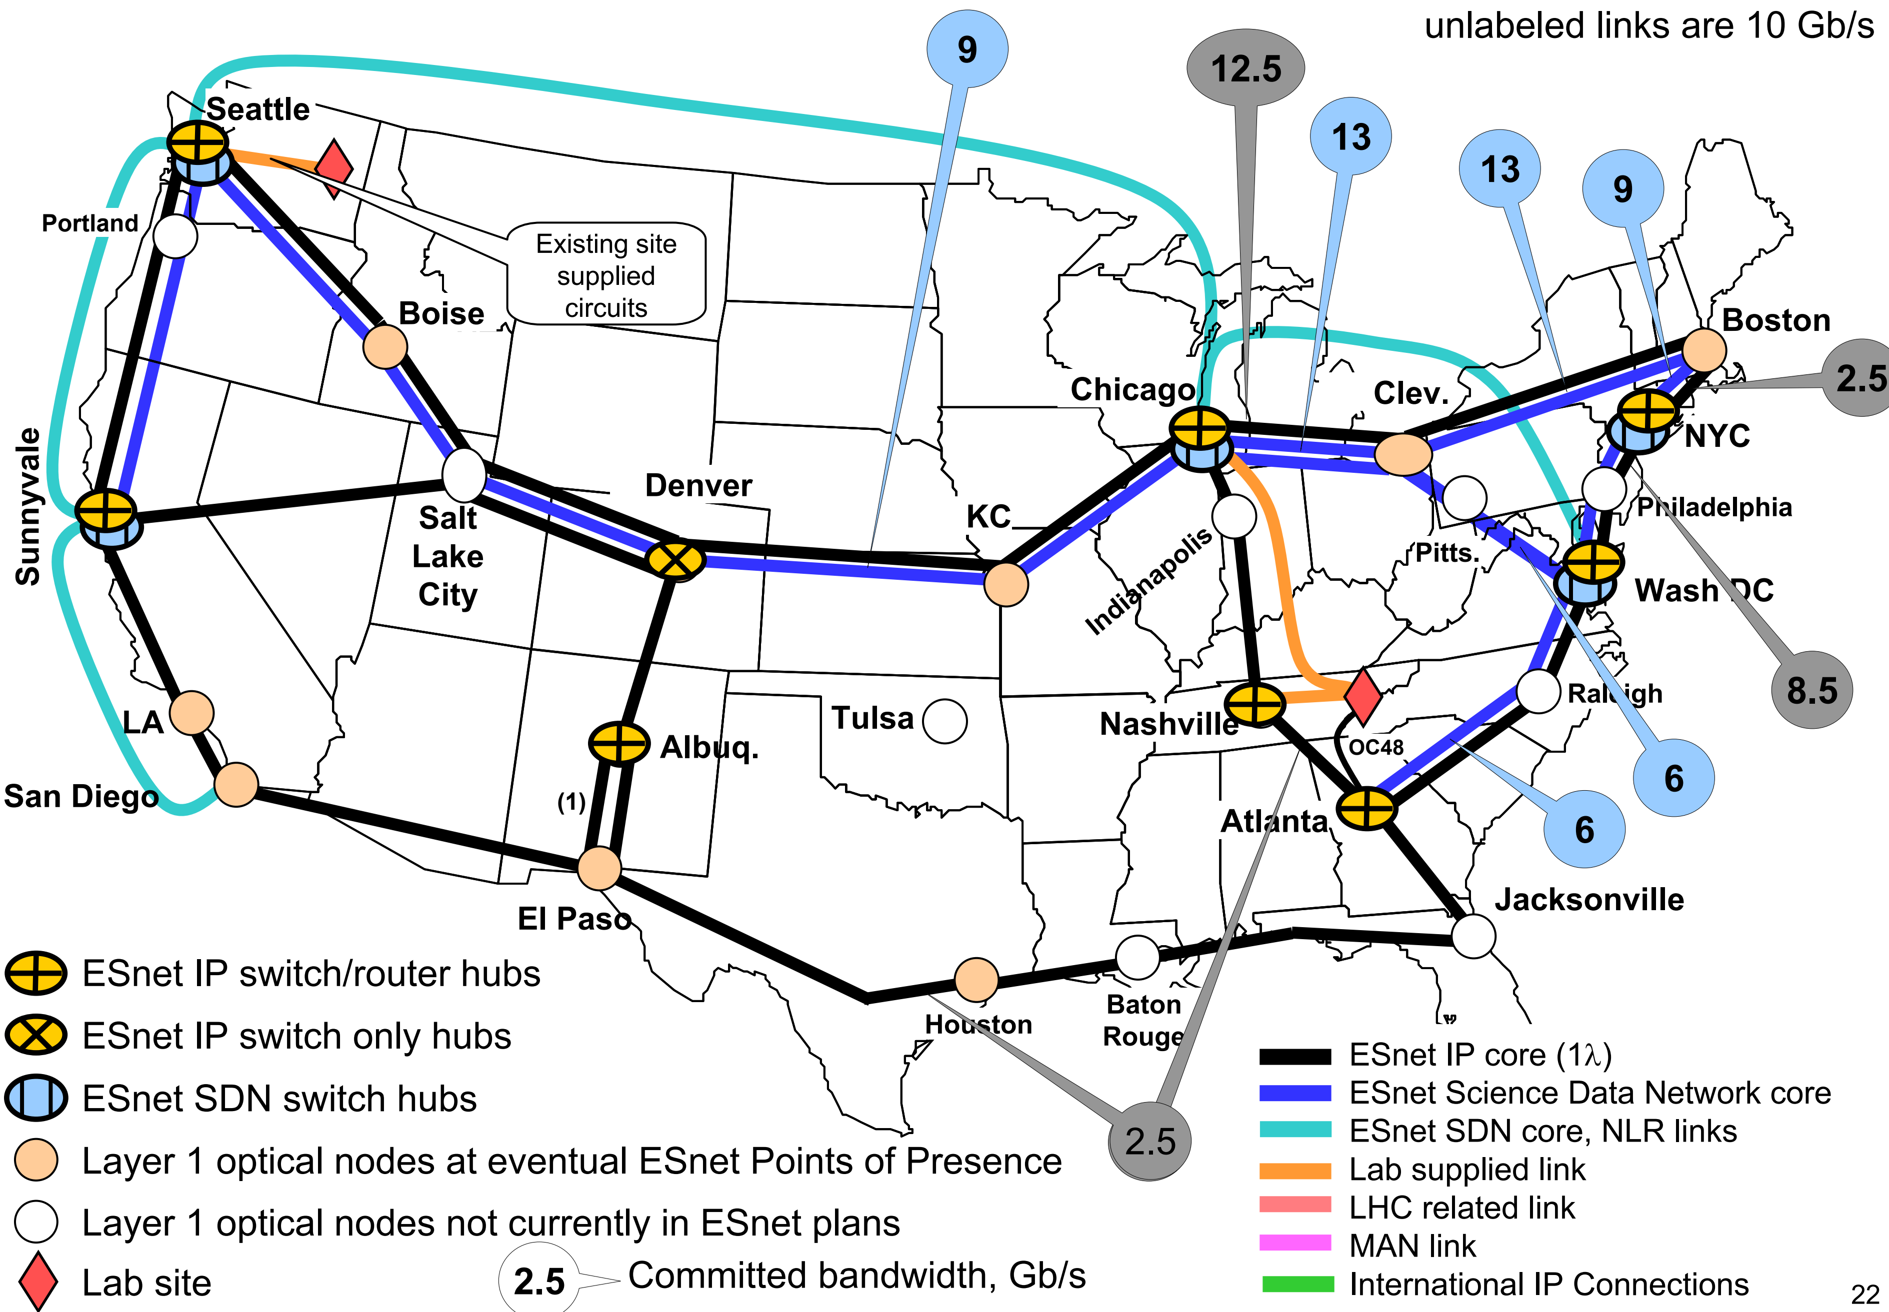

Estimated Aggregate Link Loadings, 2010-11

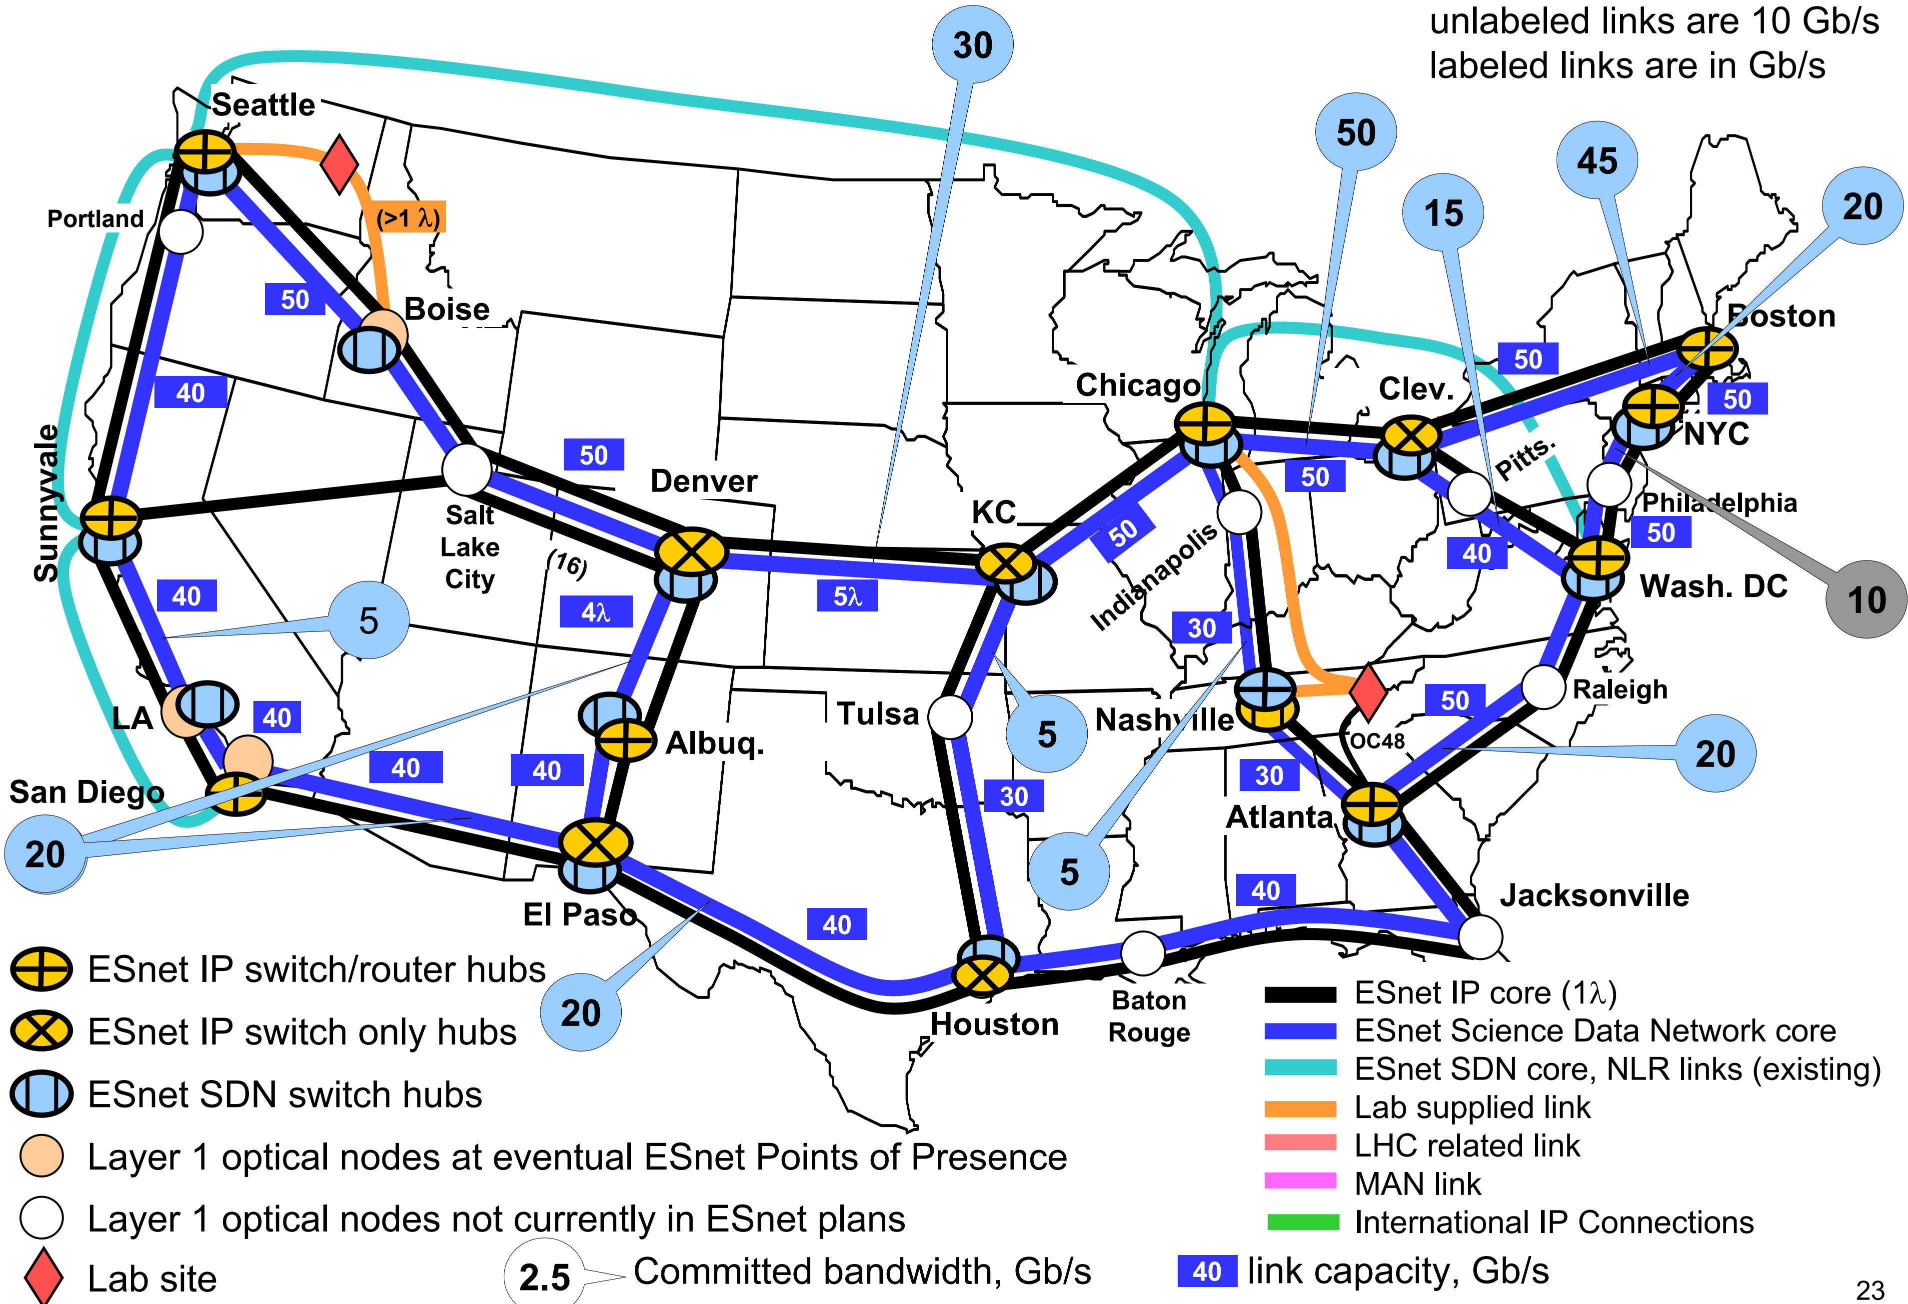

➢ESnet4 - The Response to the Requirements

I) A new network architecture and implementation strategy

- Provide two networks: IP and circuit-oriented Science Data Network
  - Reduces cost of handling high bandwidth data flows
    - Highly capable routers are not necessary when every packet goes to the same place
    - Use lower cost (factor of 5x) switches to relatively route the packets
- Rich and diverse network topology for flexible management and high reliability
- Dual connectivity at every level for all large-scale science sources and sinks
- A partnership with the US research and education community to build a shared, large-scale, R&E managed optical infrastructure
  - a scalable approach to adding bandwidth to the network
  - dynamic allocation and management of optical circuits

II) Development and deployment of a virtual circuit service

- Develop the service cooperatively with the networks that are intermediate between DOE Labs and major collaborators to ensure and-to-end interoperability

III) Develop and deploy service-oriented, user accessable network monitoring systems

IV) Provide "consulting" on system / application network performance tuning

➢ Next Generation ESnet: I) Architecture and Configuration

- **Main architectural elements and the rationale for each element**

1) A ***High-reliability IP core*** (e.g. the current ESnet core) to address
   - General science requirements
   - Lab operational requirements
   - Backup for the SDN core
   - Vehicle for science services
   - Full service IP routers

2) ***Metropolitan Area Network*** (MAN) rings to provide
   - Dual site connectivity for reliability
   - Much higher site-to-core bandwidth
   - Support for both production IP and circuit-based traffic
   - Multiply connecting the SDN and IP cores

2a) ***Loops off of the backbone*** rings to provide
   - For dual site connections where MANs are not practical

3) A ***Science Data Network*** (SDN) core for
   - Provisioned, guaranteed bandwidth circuits to support large, high-speed science data flows
   - Very high total bandwidth
   - Multiply connecting MAN rings for protection against hub failure
   - Alternate path for production IP traffic
   - Less expensive router/switches
   - Initial configuration targeted at LHC, which is also the first step to the general configuration that will address all SC requirements
   - Can meet other unknown bandwidth requirements by adding lambdas

ESnet Target Architecture:
IP WAN+Science Data Network WAN+Metro Area Rings

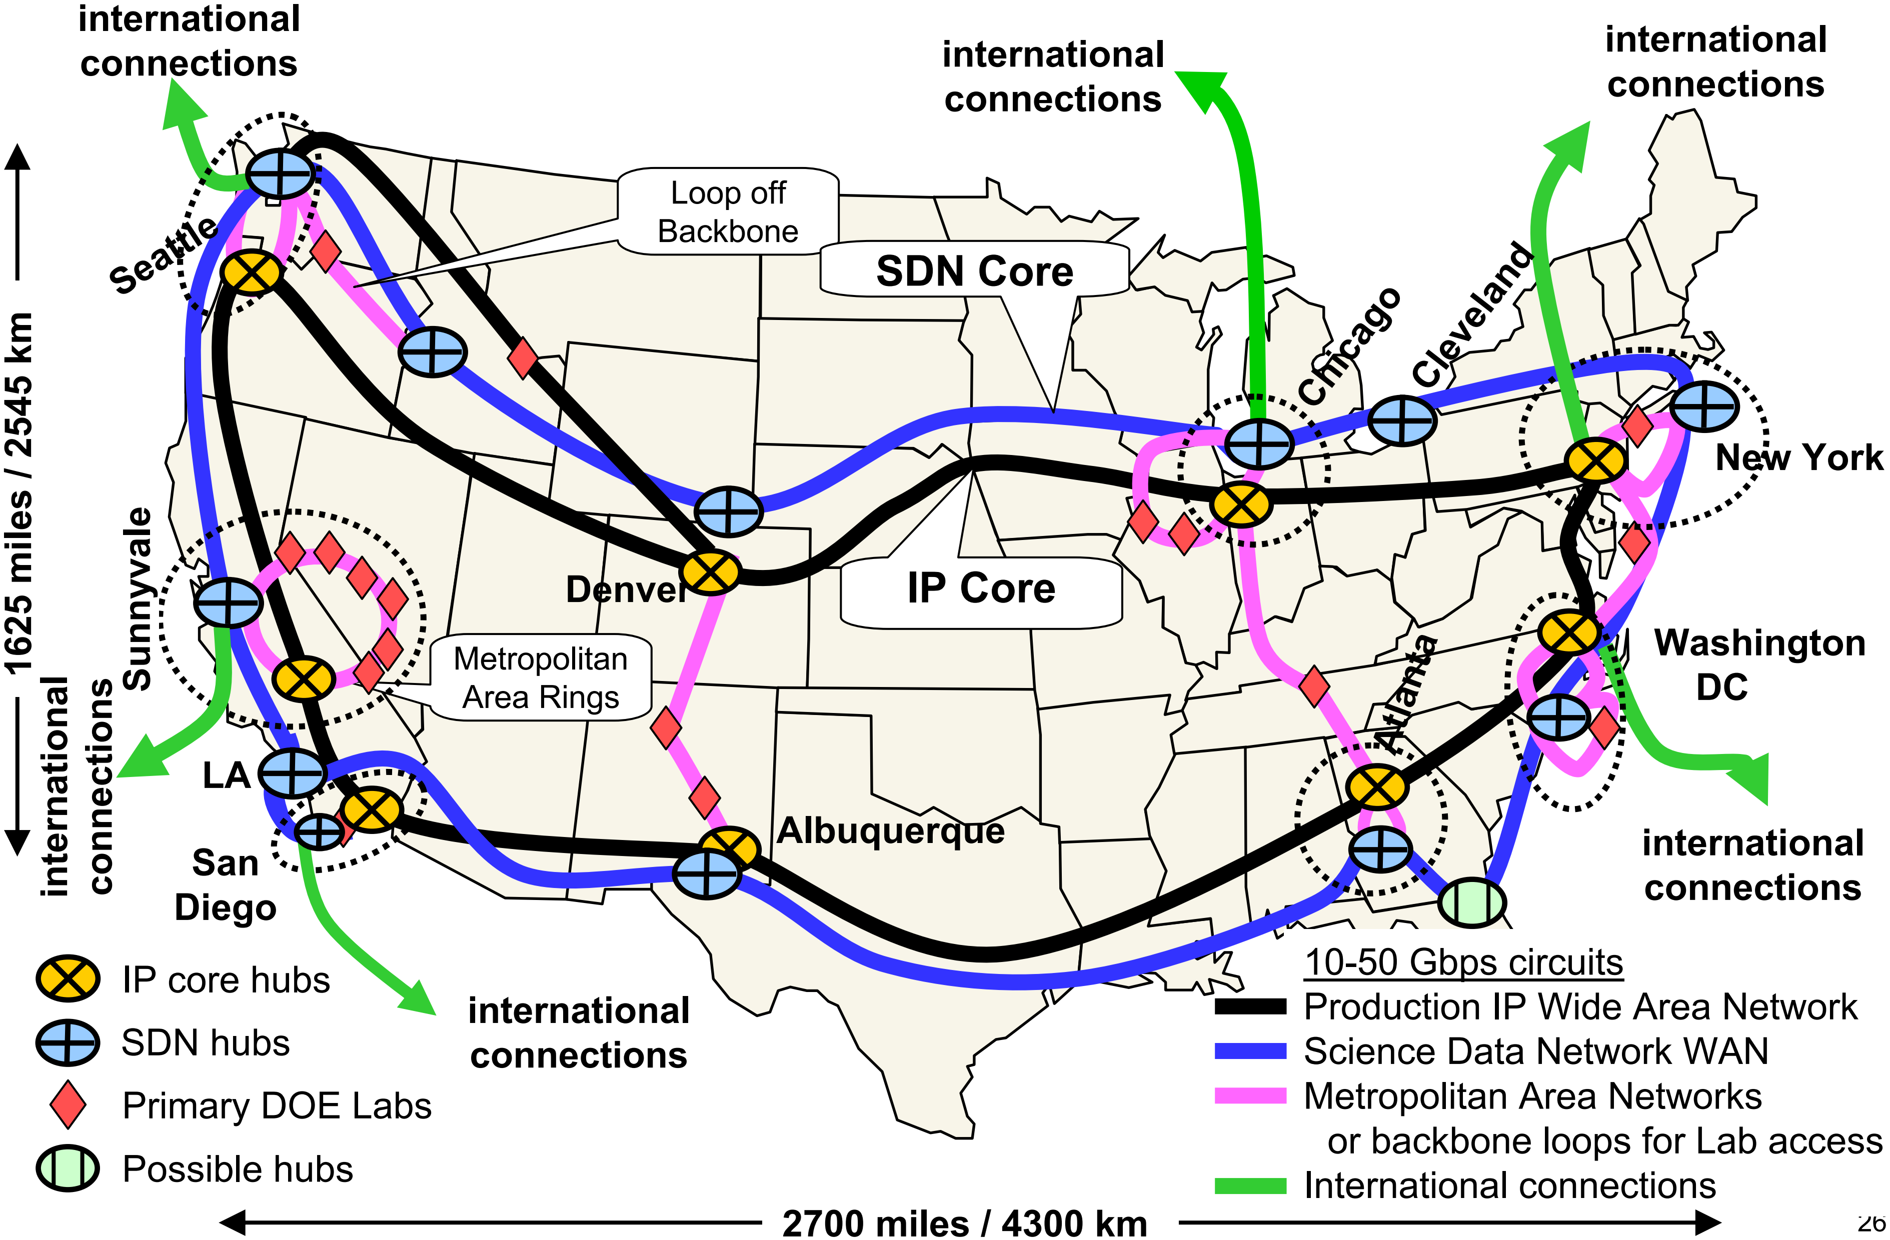

ESnet Metropolitan Area Network Ring Architecture for High Reliability Sites

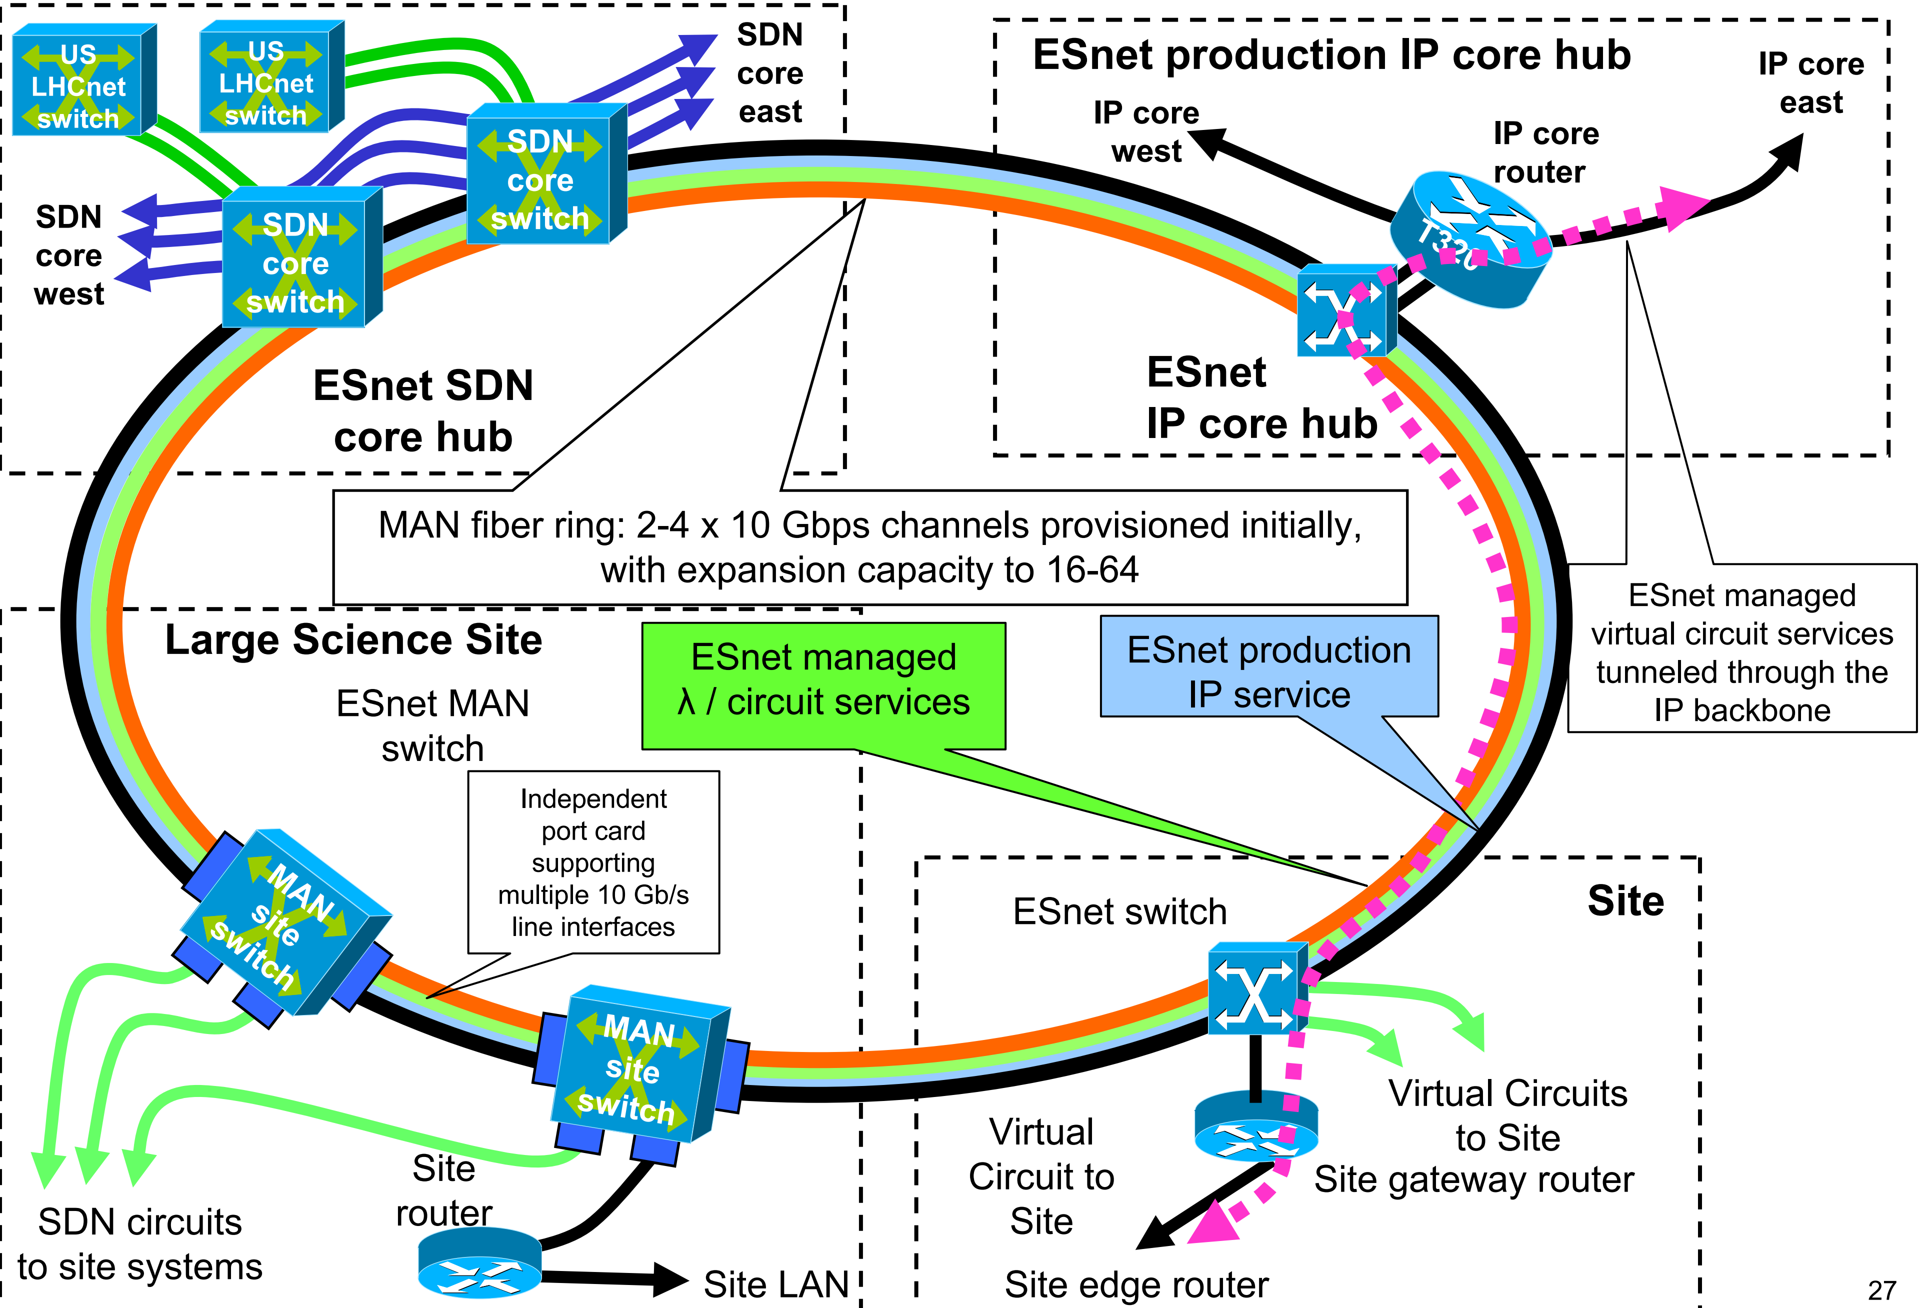

ESnet4

- Internet2 has partnered with Level 3 Communications Co. and Infinera Corp. for a dedicated optical fiber infrastructure with a national footprint and a rich topology - the “Internet2 Network”
  - The fiber will be provisioned with Infinera Dense Wave Division Multiplexing equipment that uses an advanced, integrated optical-electrical design
  - Level 3 will maintain the fiber and the DWDM equipment
  - The DWDM equipment will initially be provisioned to provide10 optical circuits (lambdas - λs) across the entire fiber footprint (40/80 λs is max.)
- ESnet has partnered with Internet2 to:
  - Share the optical infrastructure
  - Develop new circuit-oriented network services
  - Explore mechanisms that could be used for the ESnet Network Operations Center (NOC) and the Internet2/Indiana University NOC to back each other up for disaster recovery purposes

ESnet4

- ESnet is building its next generation IP network and its new circuit-oriented Science Data Network primarily on the Internet2 circuits (λs) that are dedicated to ESnet, together with a few National Lambda Rail and other circuits
  - ESnet will provision and operate its own routing and switching hardware that is installed in various commercial telecom hubs around the country, as it has done for the past 20 years
  - ESnet's peering relationships with the commercial Internet, various US research and education networks, and numerous international networks will continue and evolve as they have for the past 20 years

ESnet4

- ESnet4 will also involve an expansion of the multi-10Gb/s Metropolitan Area Rings in the San Francisco Bay Area, Chicago, Long Island, Newport News (VA/Washington, DC area), and Atlanta
  - provide multiple, independent connections for ESnet sites to the ESnet core network
  - expandable
- Several 10Gb/s links provided by the Labs that will be used to establish multiple, independent connections to the ESnet core
  - currently PNNL and ORNL

Internet2 and ESnet Optical Node

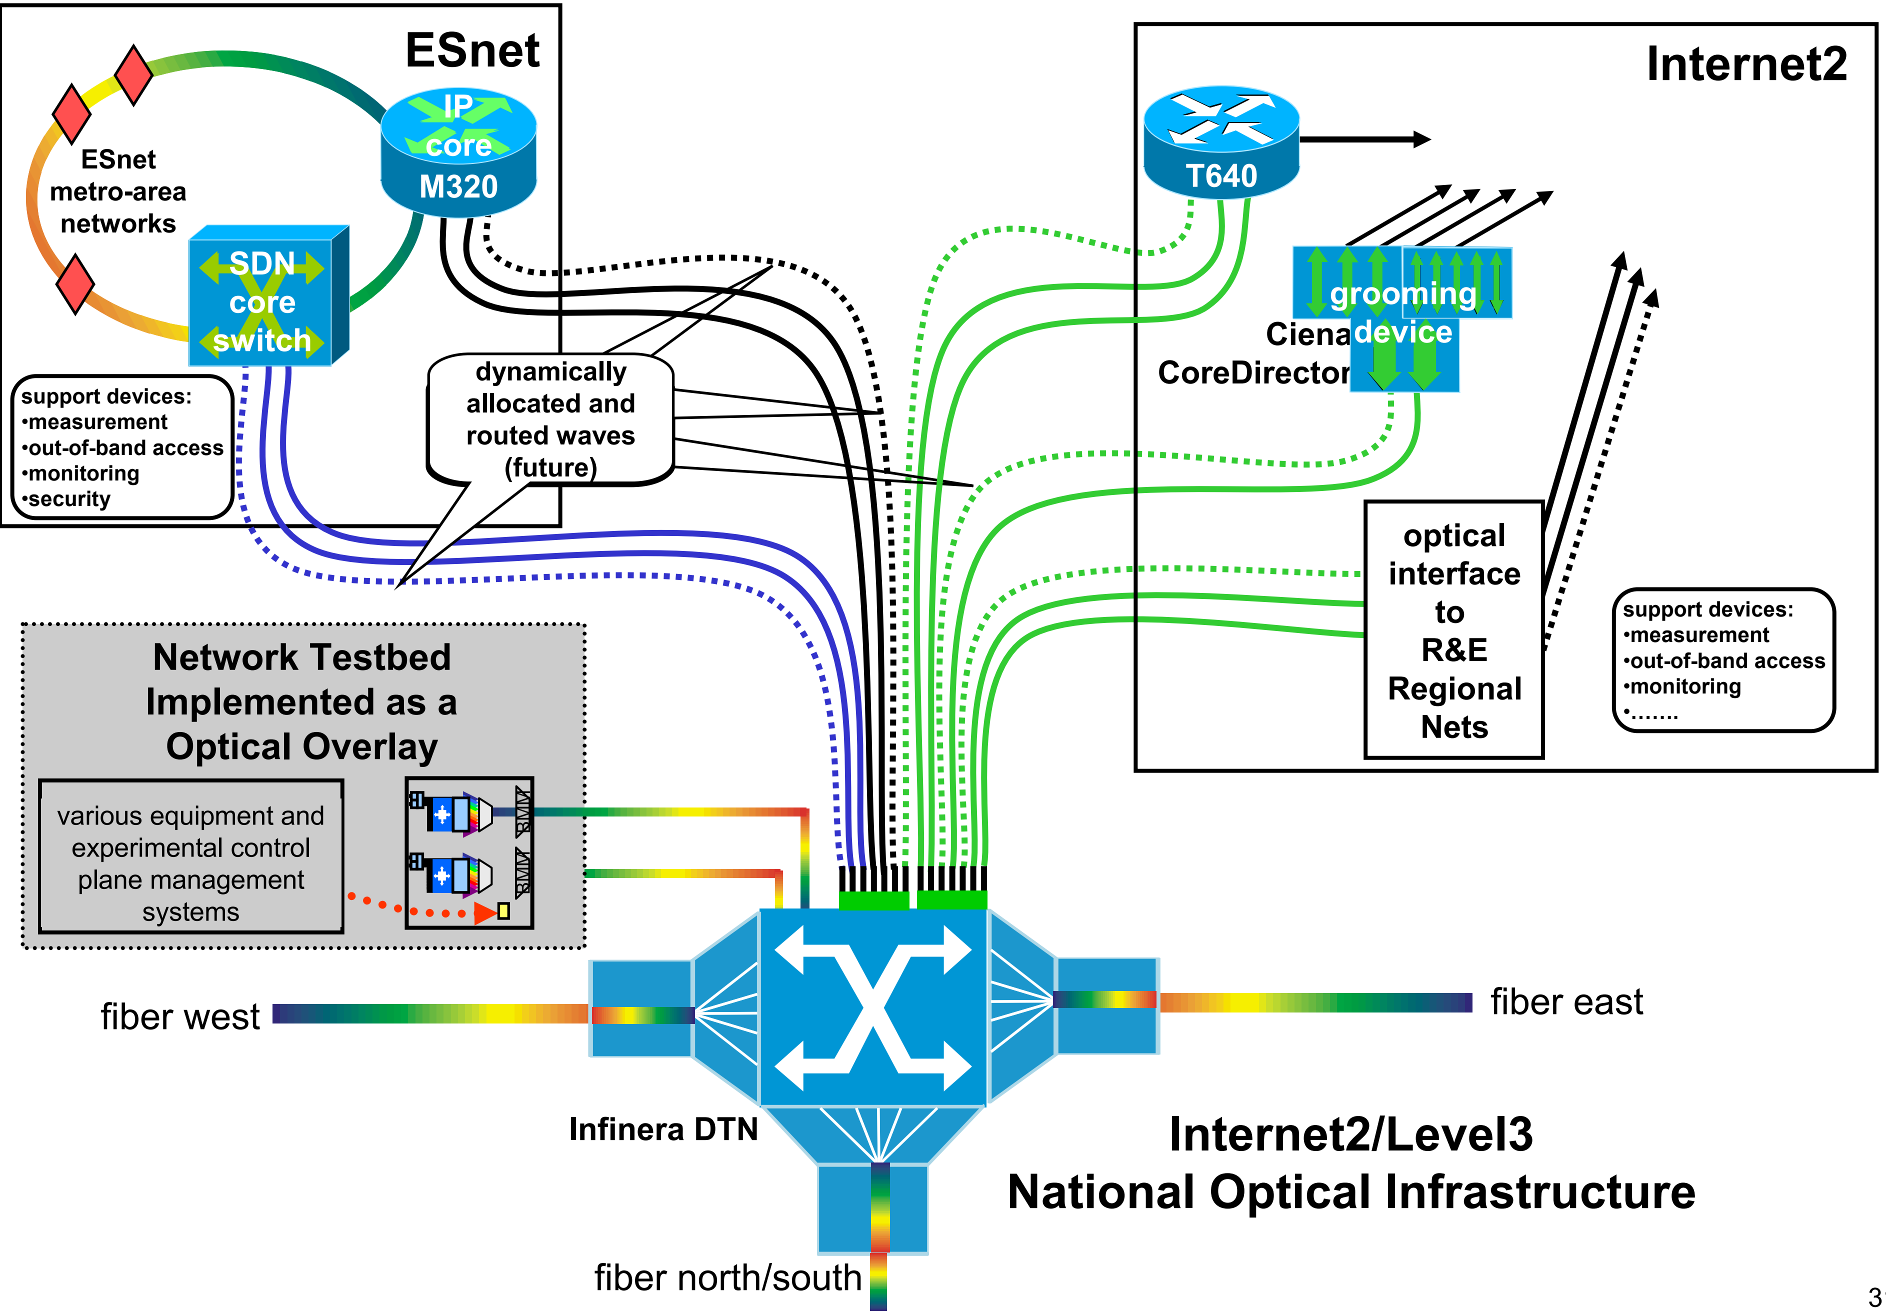

ESnet 3 Backbone as of January 1, 2007

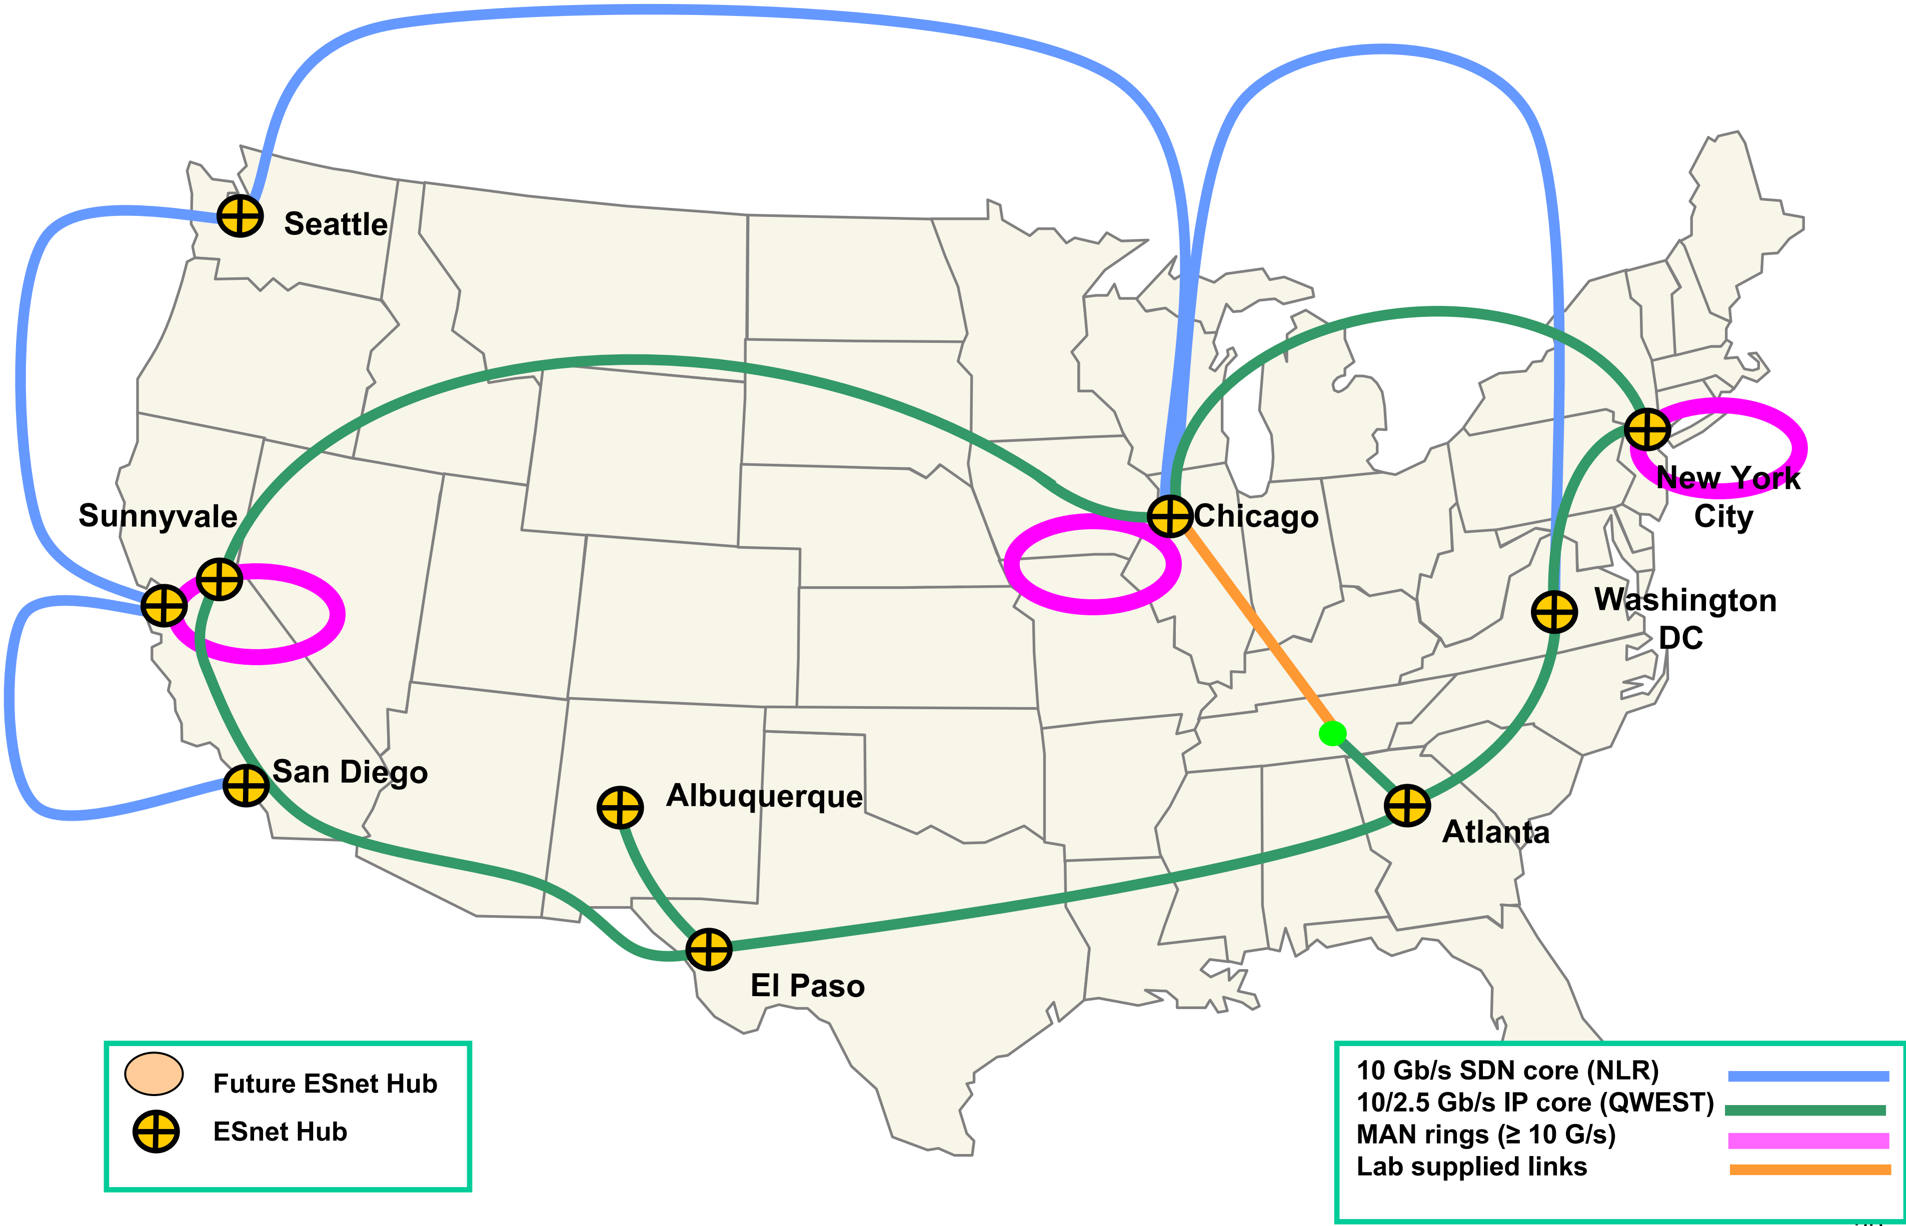

ESnet 4 Backbone Target September 30, 2007

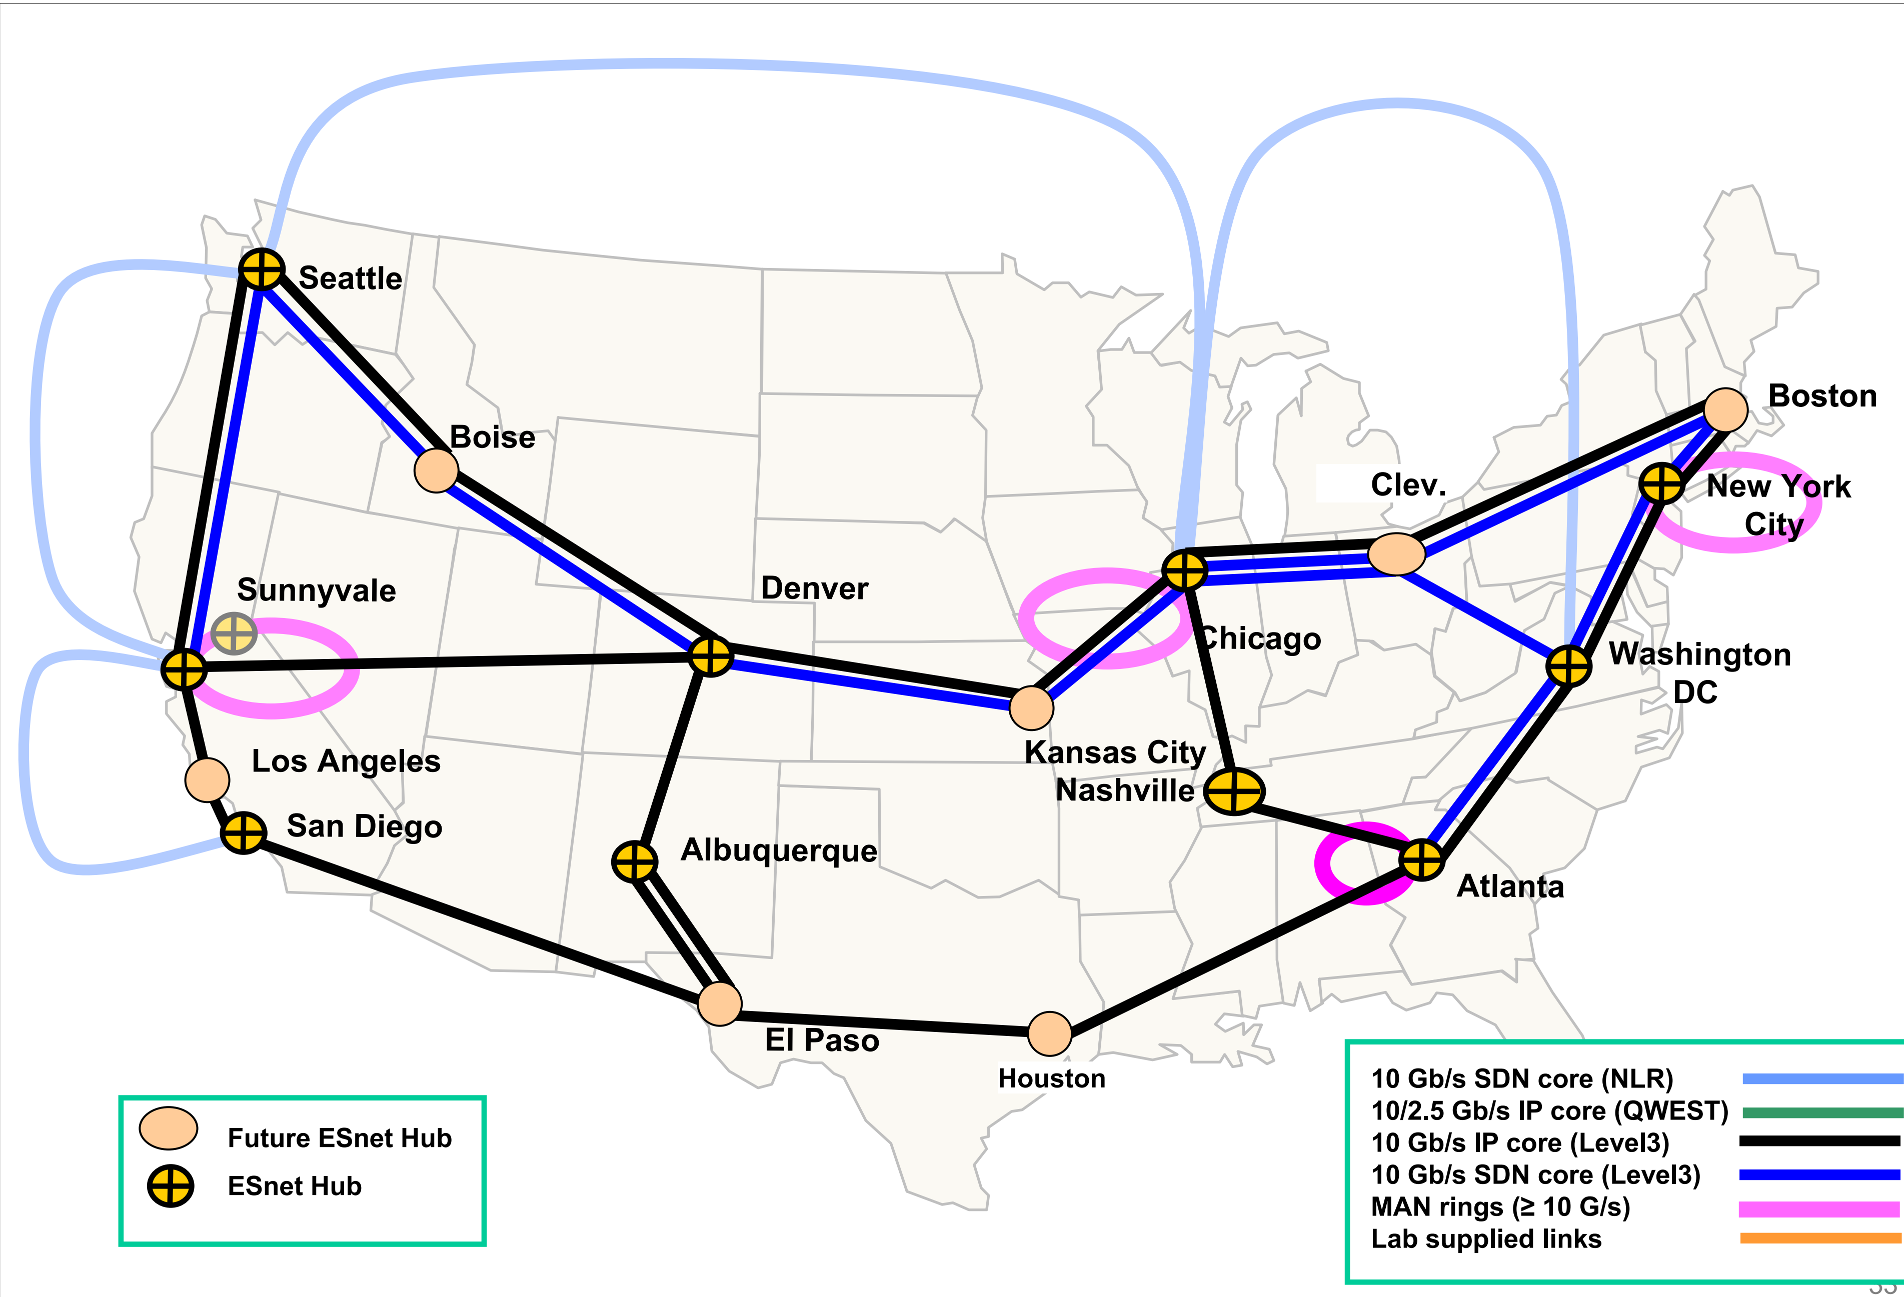

ESnet4 Metro Area Rings, 2007 Configurations

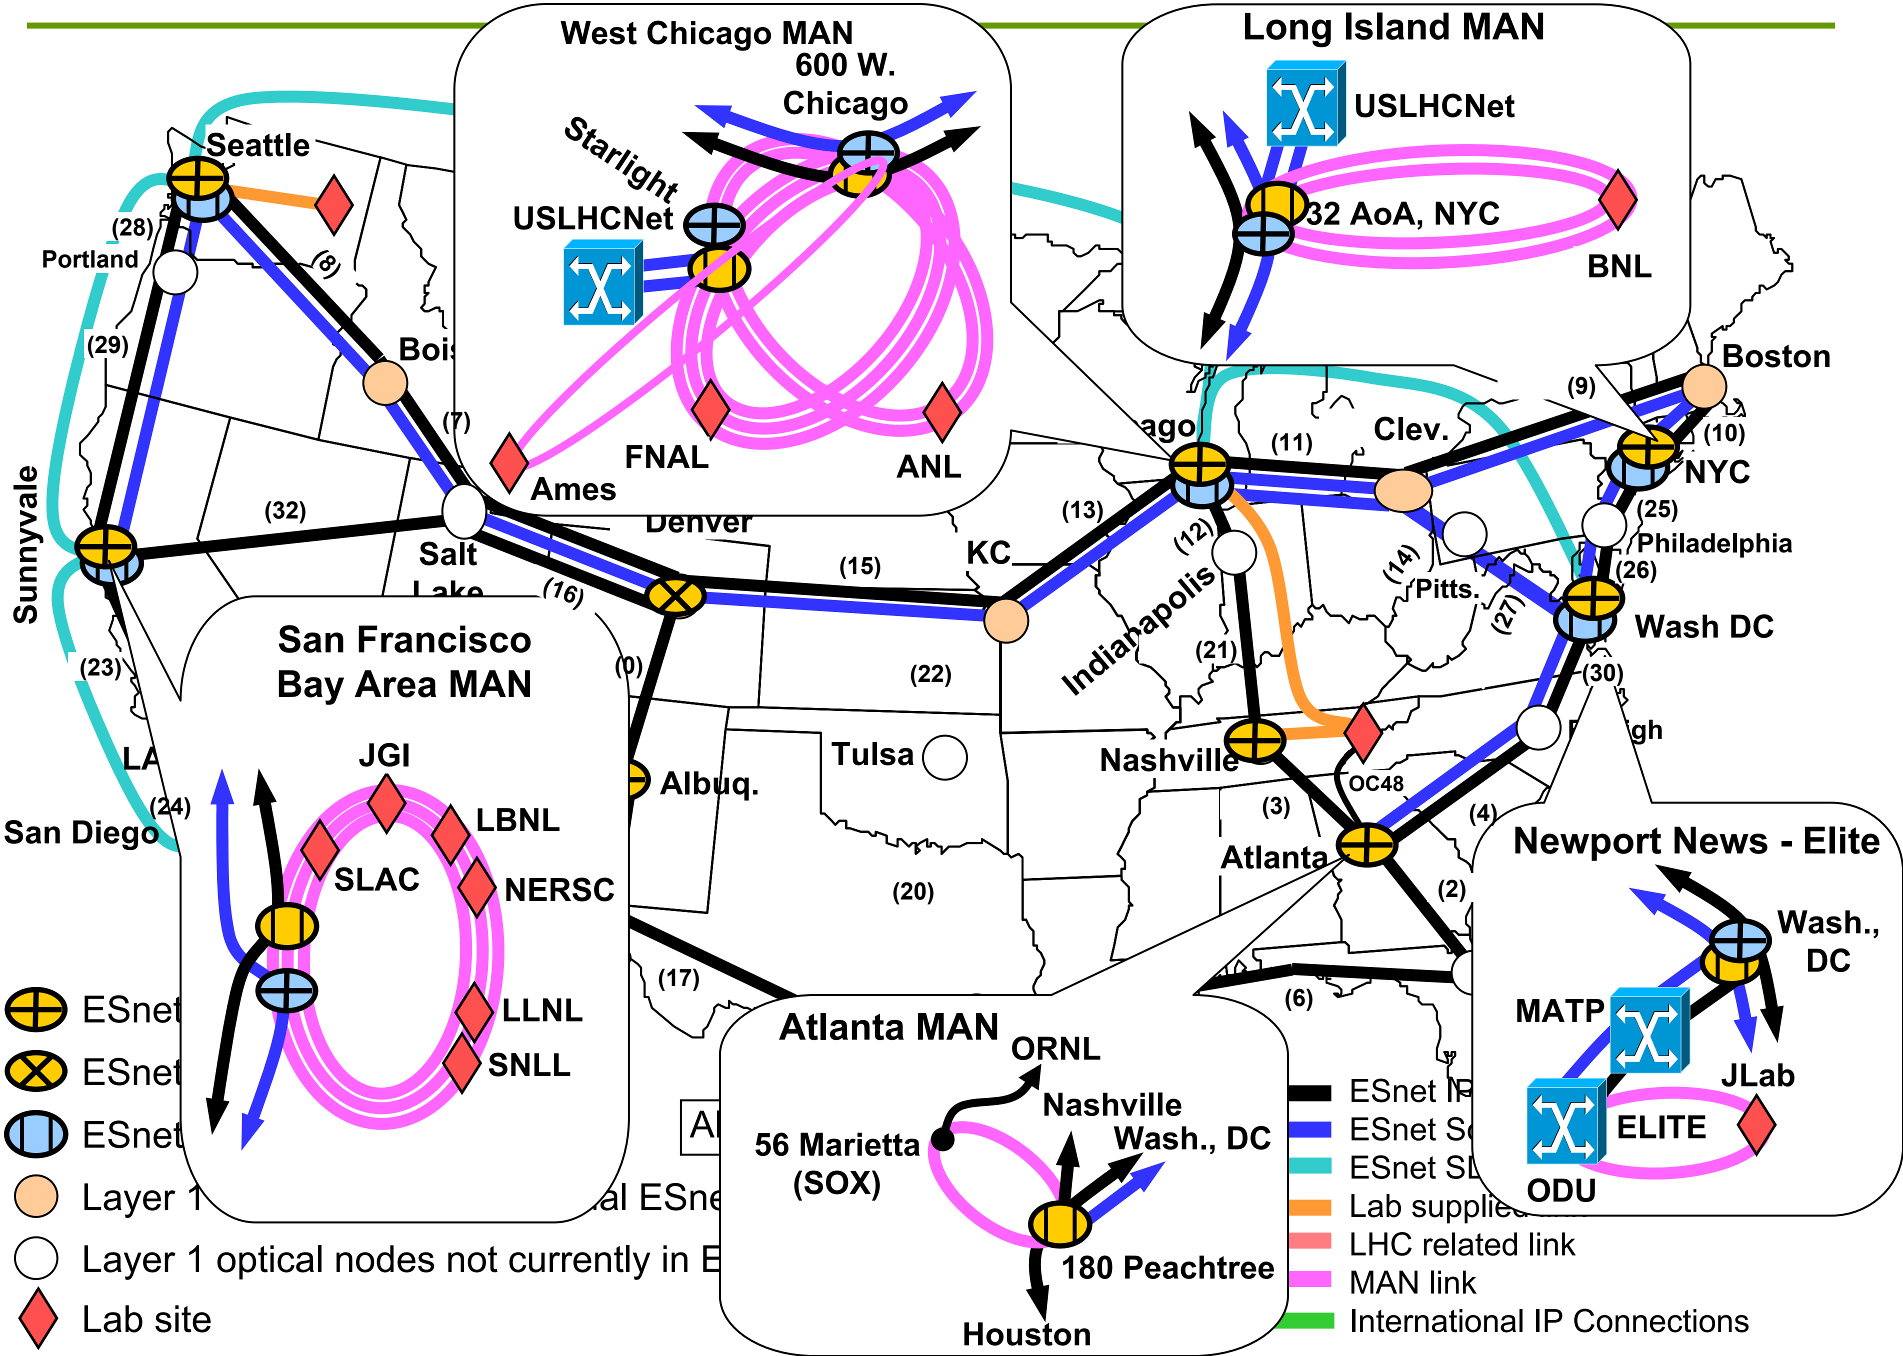

ESnet4 IP + SDN, 2008 Configuration (Estimated)

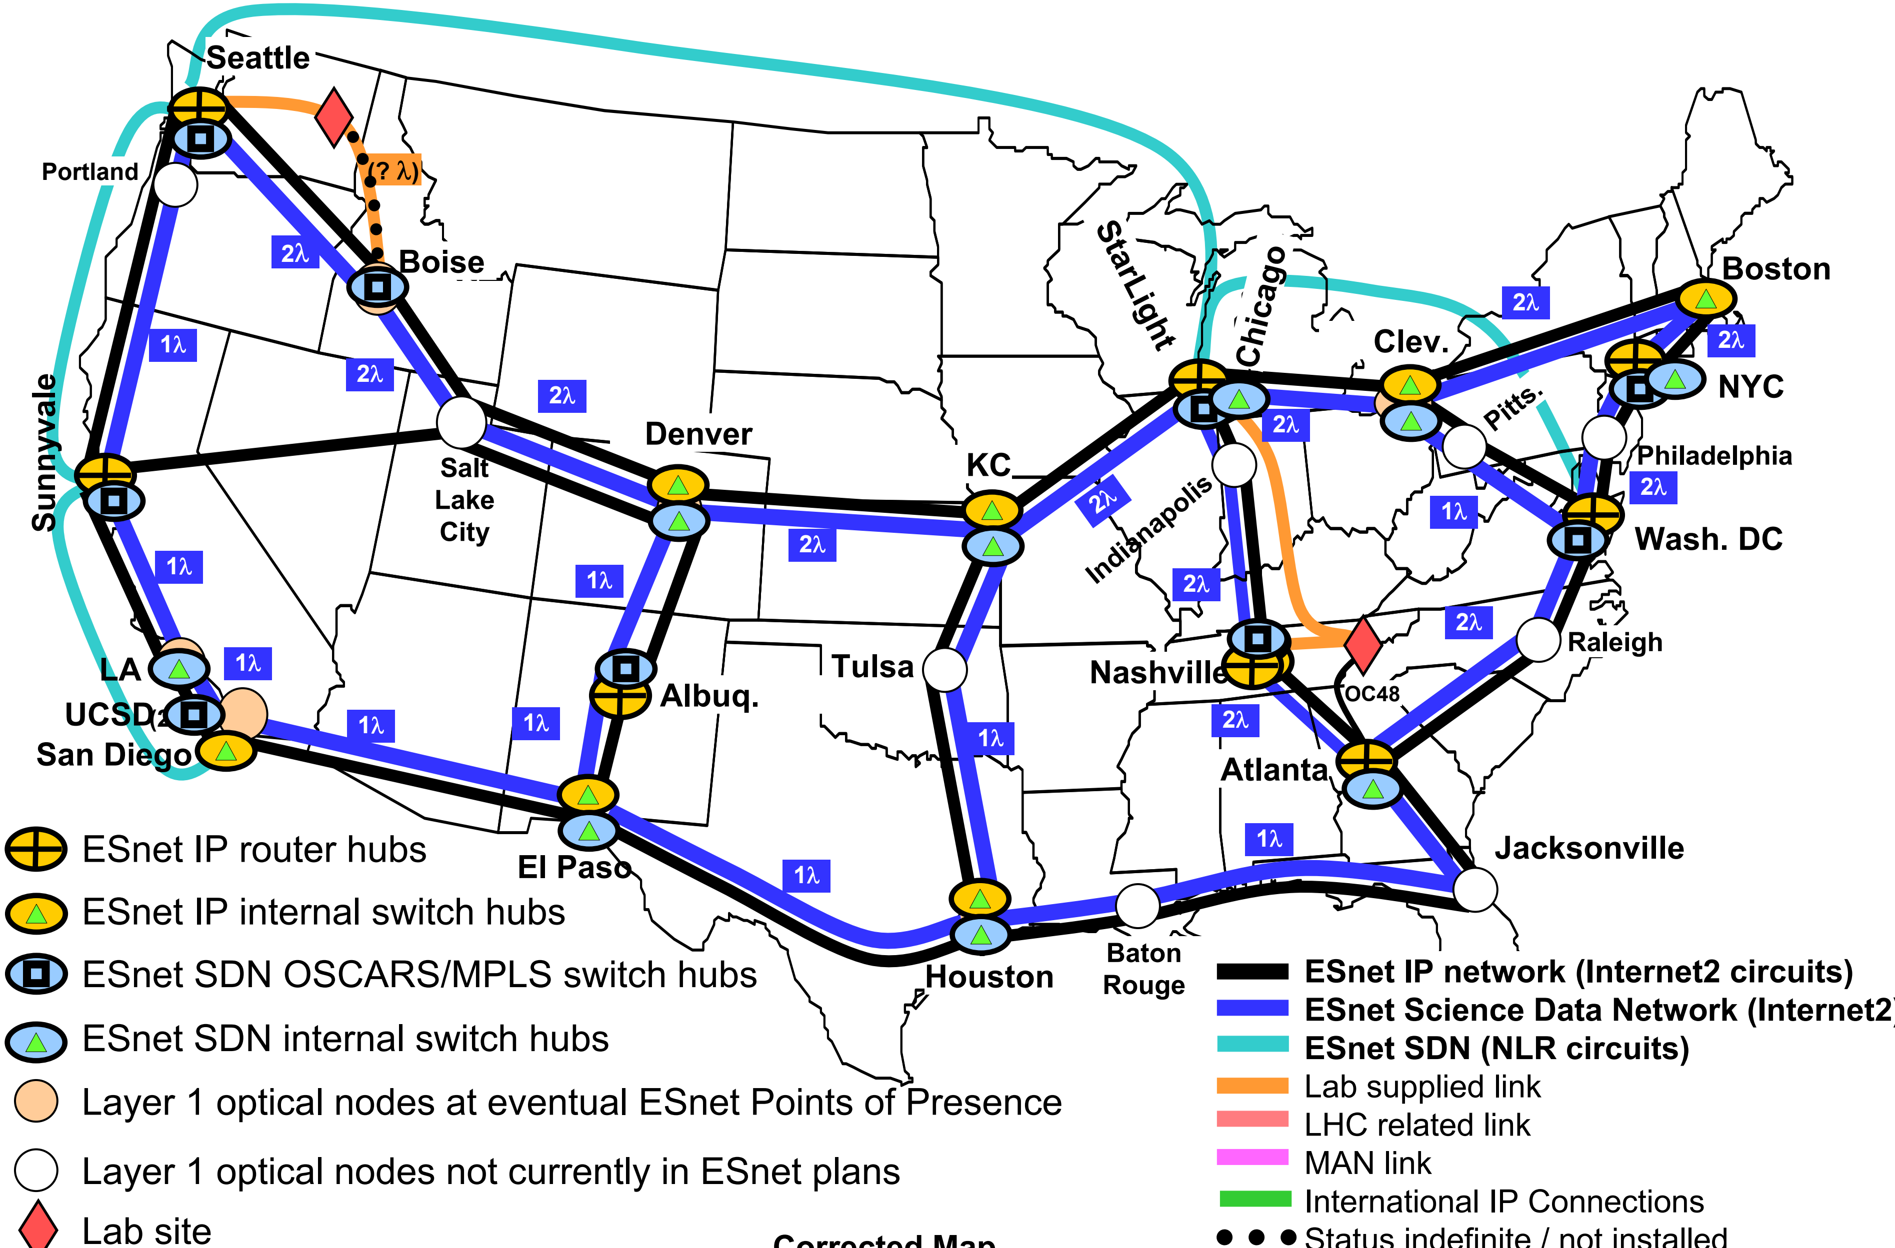

Corrected Map

Estimated ESnet4 2009 Configuration
(Some of the circuits may be allocated dynamically from shared a pool.)

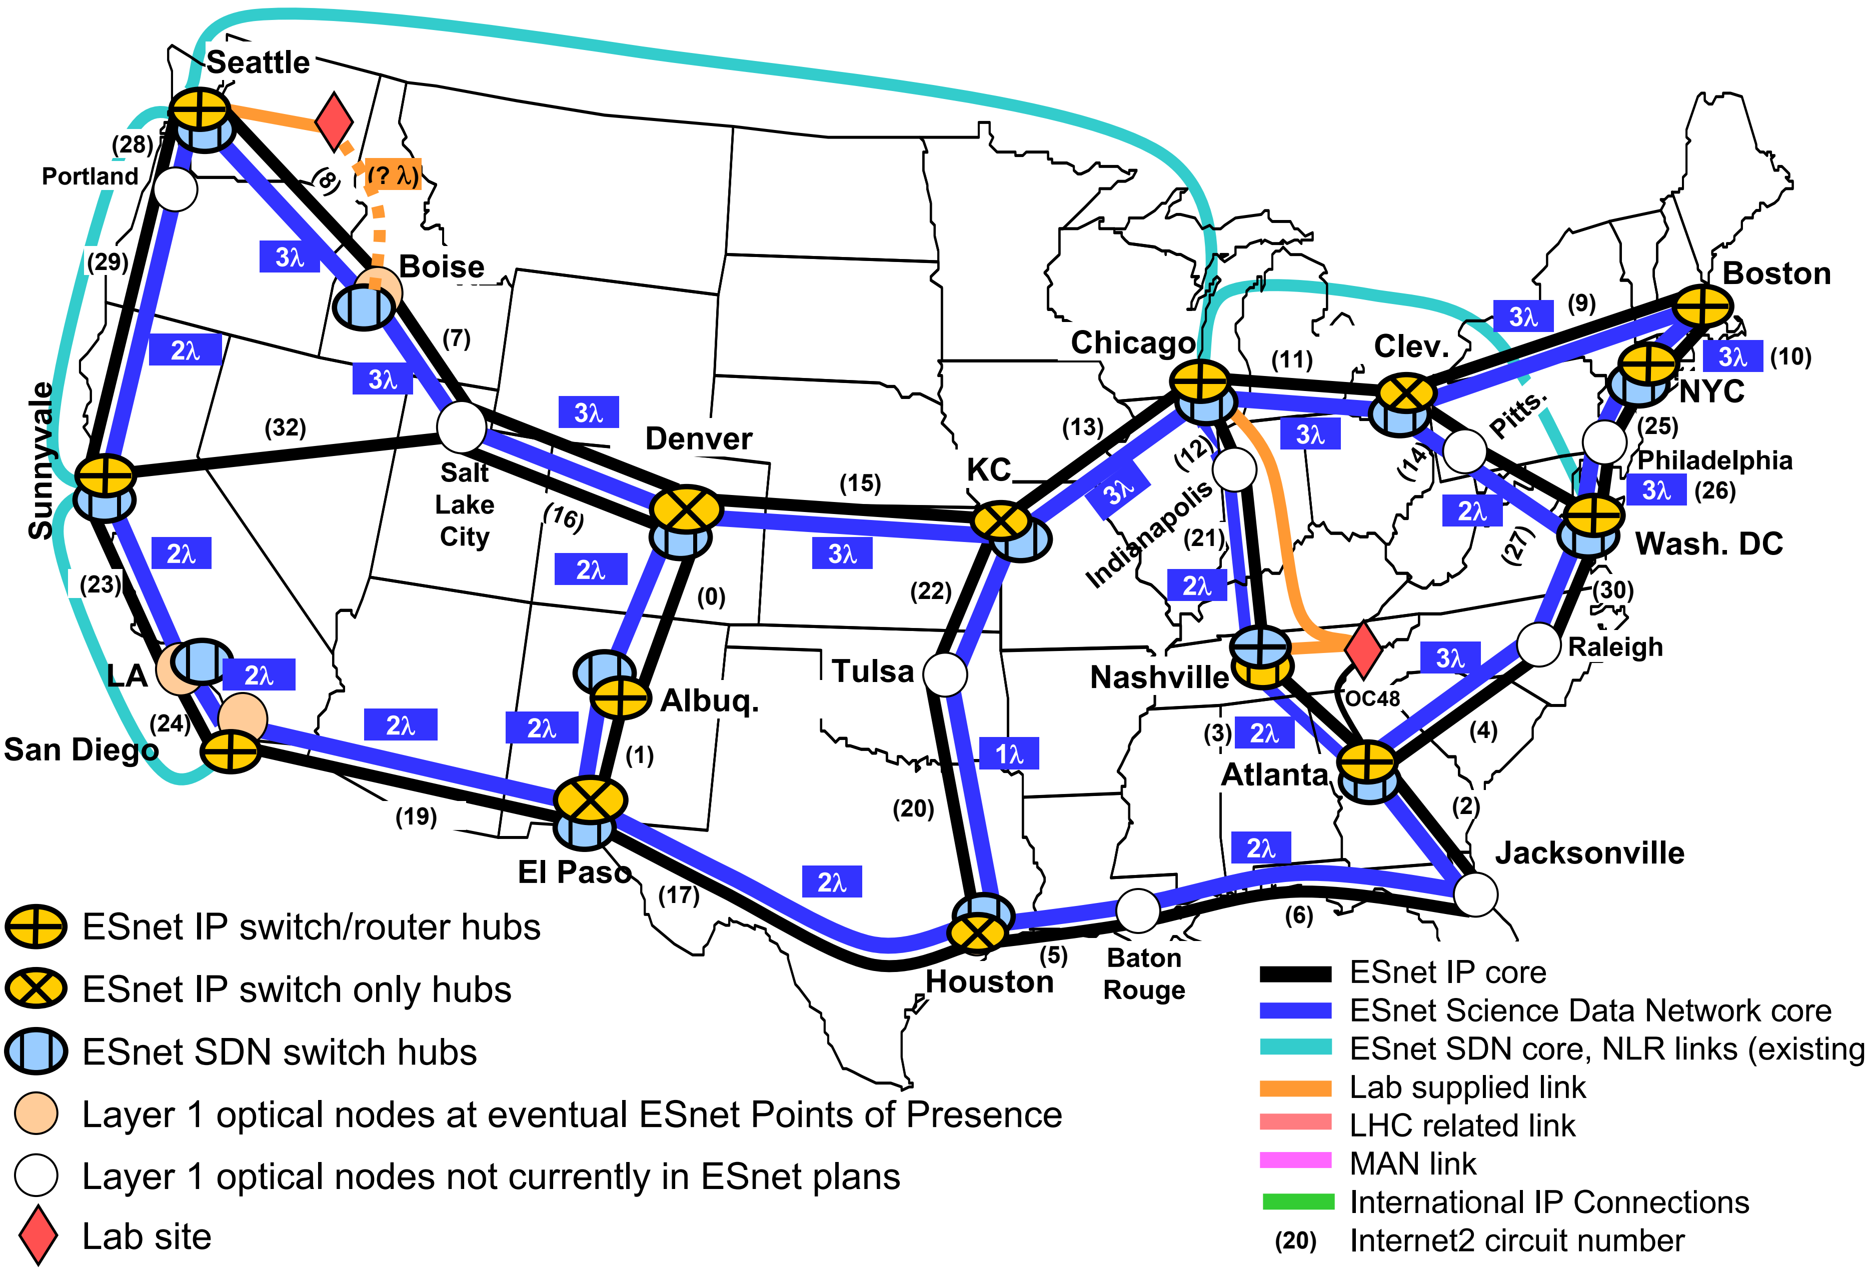

ESnet4 IP + SDN, 2011 Configuration

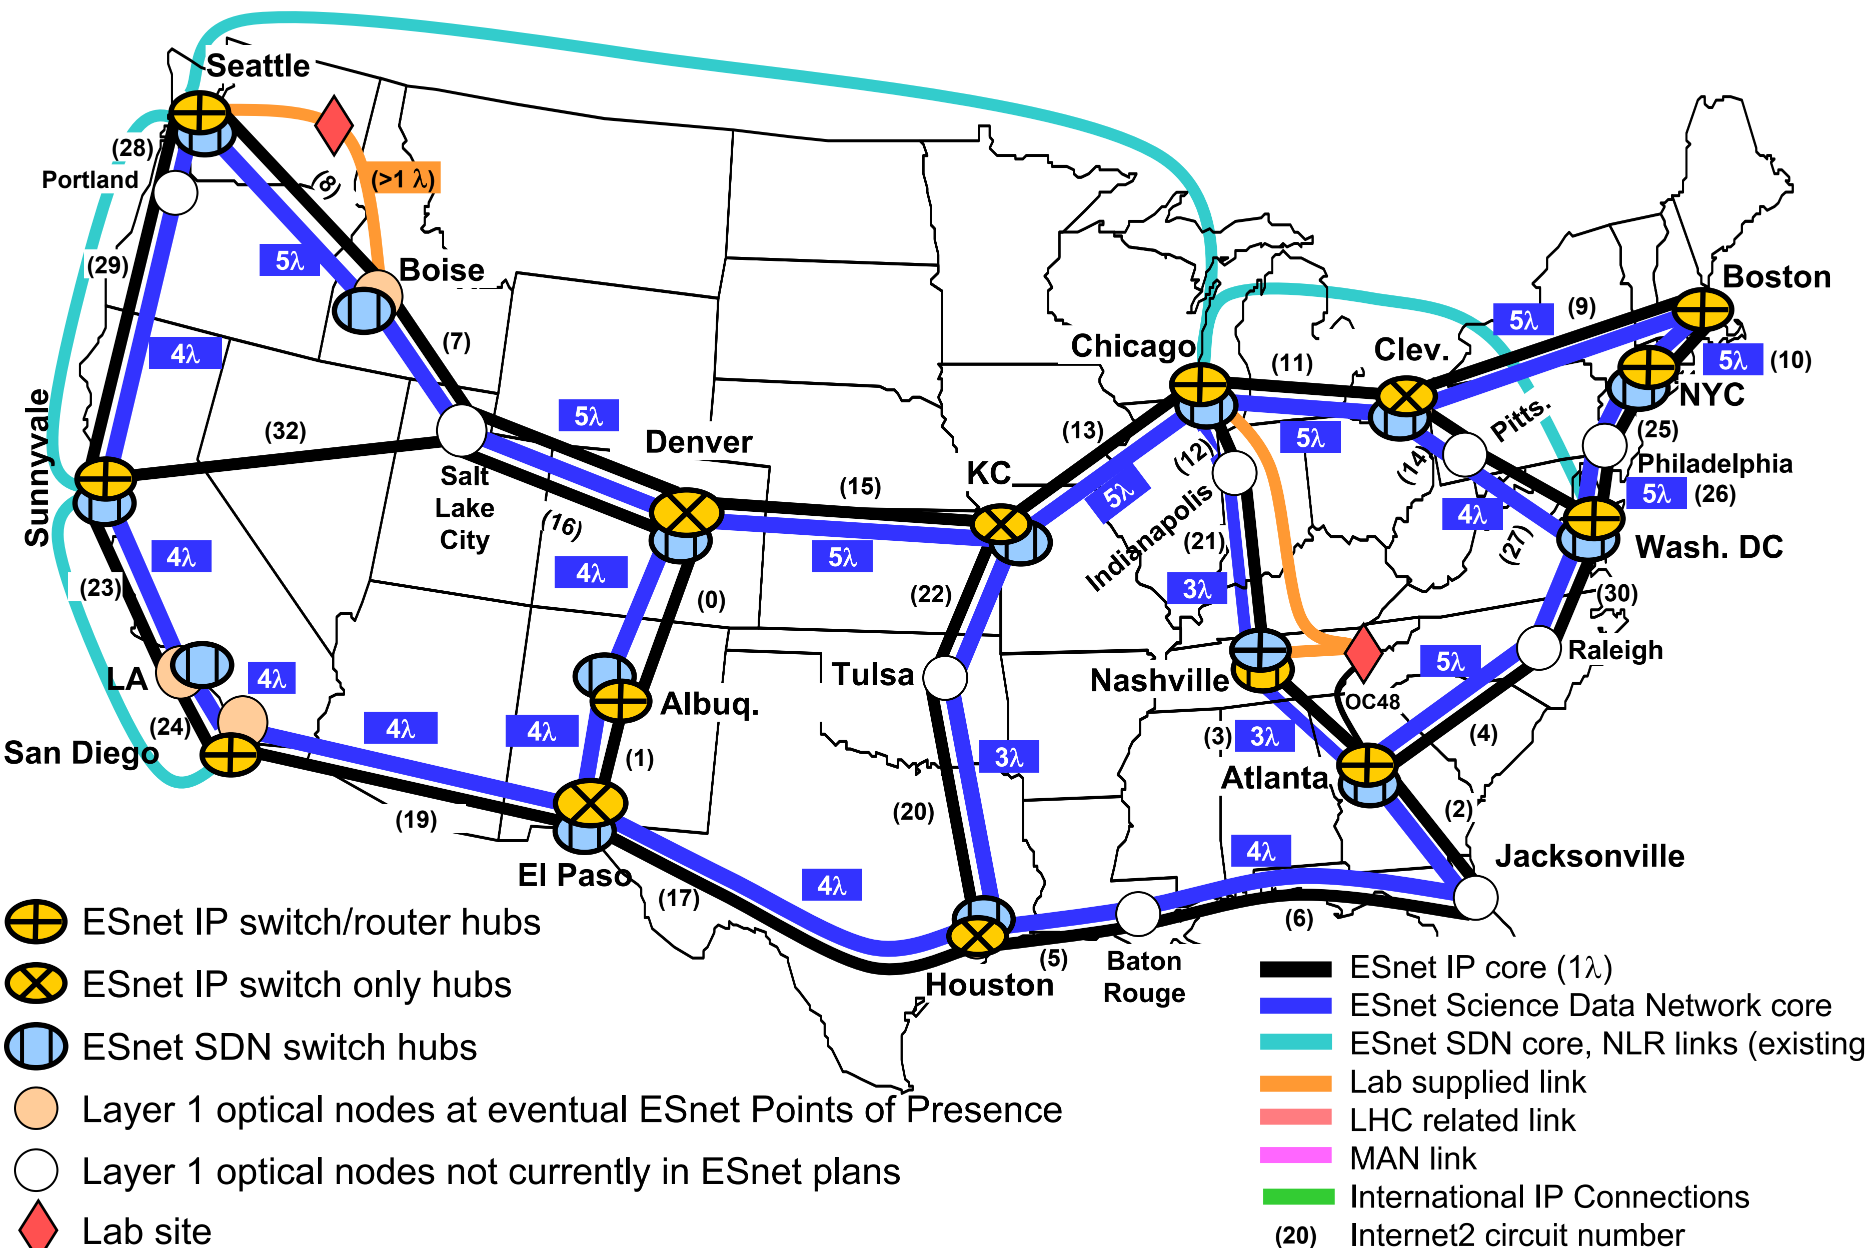

ESnet4 Planed Configuration

40-50 Gbps in 2009-2010, 160-400 Gbps in 2011-2012

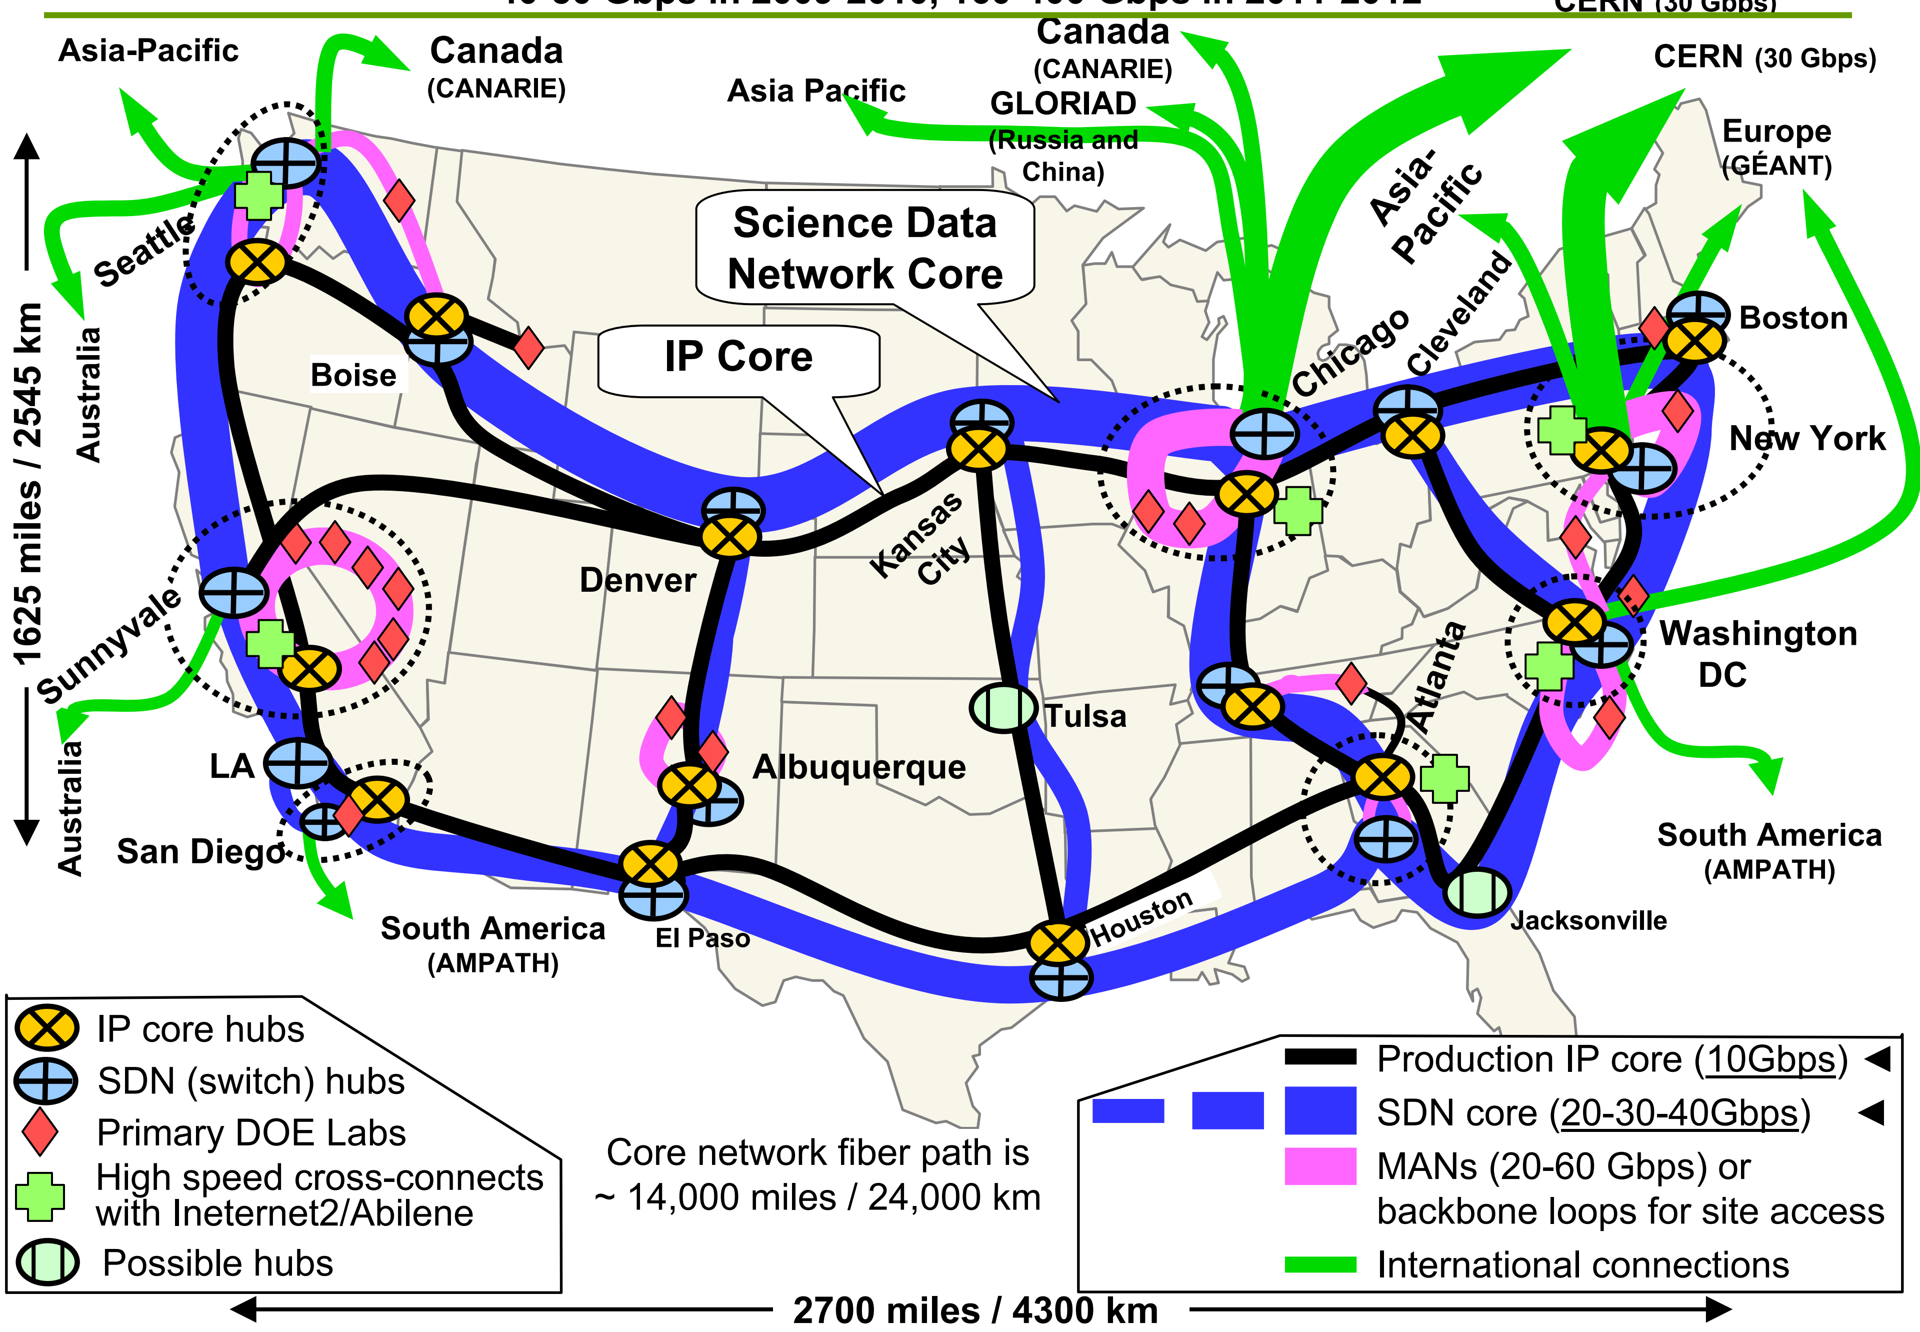

New Network Service: Virtual Circuits

- Guaranteed bandwidth service
  - User specified bandwidth - requested and managed in a Web Services framework
- Traffic isolation and traffic engineering
  - Provides for high-performance, non-standard transport mechanisms that cannot co-exist with commodity TCP-based transport
  - Enables the engineering of explicit paths to meet specific requirements
    - e.g. bypass congested links, using lower bandwidth, lower latency paths
- End-to-end (cross-domain) connections between Labs and collaborating institutions
- Secure connections
  - Provides end-to-end connections between Labs and collaborator institutions
  - The circuits are “secure” to the edges of the network (the site boundary) because they are managed by the control plane of the network which is isolated from the general traffic
- Reduced cost of handling high bandwidth data flows
  - Highly capable routers are not necessary when every packet goes to the same place
  - Use lower cost (factor of 5x) switches to relatively route the packets

Virtual Circuit Service Functional Requirements

- Support user/application VC reservation requests
  - Source and destination of the VC
  - Bandwidth, latency, start time, and duration of the VC
  - Traffic characteristics (e.g. flow specs) to identify traffic designated for the VC
- Manage allocations of scarce, shared resources
  - Authentication to prevent unauthorized access to this service
  - Authorization to enforce policy on reservation/provisioning
  - Gathering of usage data for accounting
- Provide virtual circuit setup and teardown mechanisms and security
  - Widely adopted and standard protocols (such as MPLS and GMPLS) are well understood within a single domain
  - Cross domain interoperability is the subject of ongoing, collaborative development
  - secure and-to-end connection setup is provided by the network control plane
  - accommodate heterogeneous circuit abstraction (e.g.. MPLS, GMPLS, VLANs, VCAT/LCAS)
- Enable the claiming of reservations
  - Traffic destined for the VC must be differentiated from “regular” traffic
- Enforce usage limits
  - Per VC admission control polices usage, which in turn facilitates guaranteed bandwidth
  - Consistent per-hop QoS throughout the network for transport predictability

Environment of Science is Inherently Multi-Domain

- End points will be at independent institutions – campuses or research institutes - that are served by ESnet, Abilene, GÉANT, and their regional networks
  - Complex inter-domain issues – typical circuit will involve five or more domains - of necessity this involves collaboration with other networks
  - For example, a connection between FNAL and DESY involves five domains, traverses four countries, and crosses seven time zones

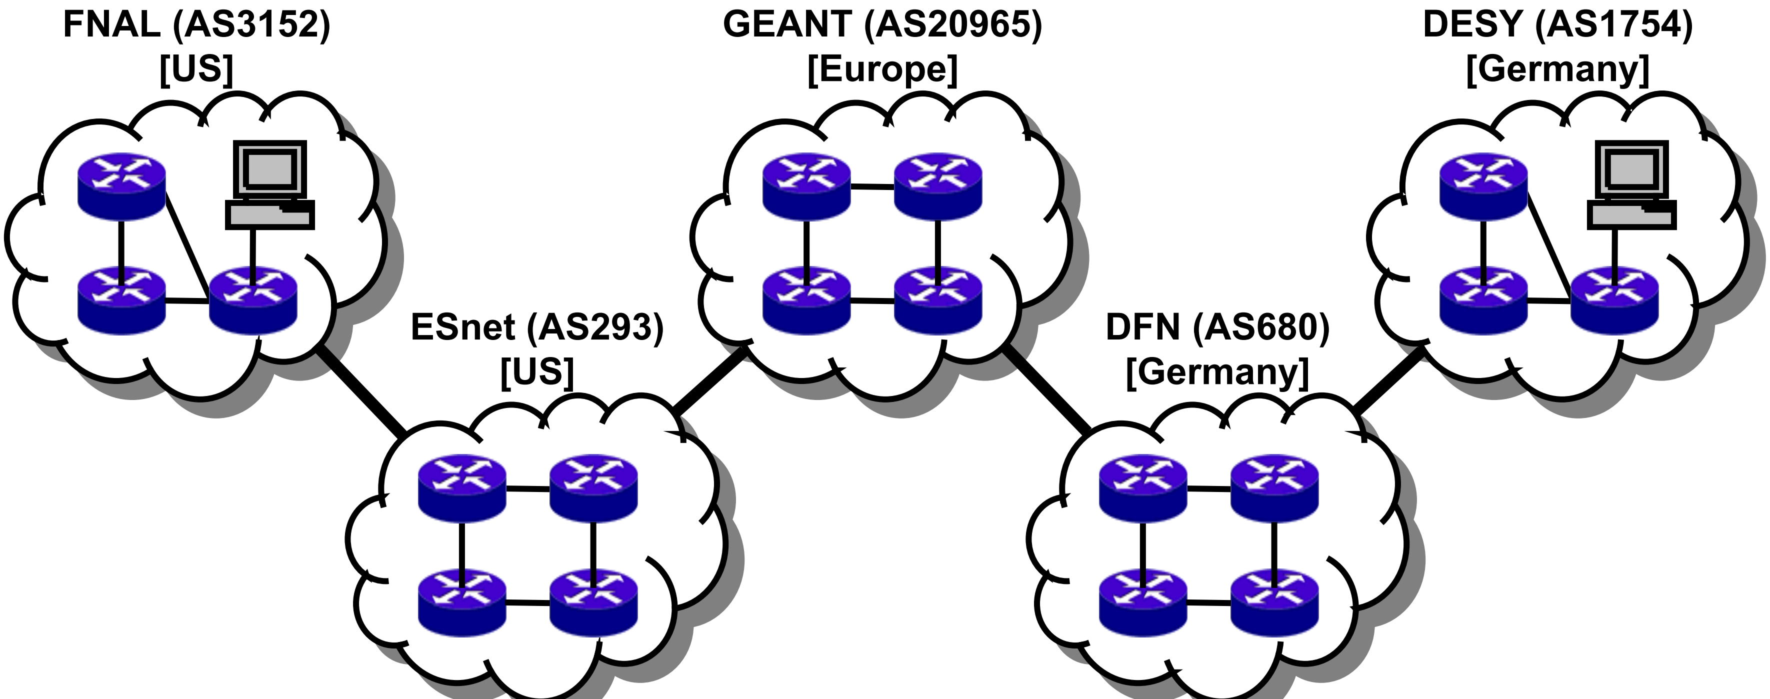

OSCARS [6]

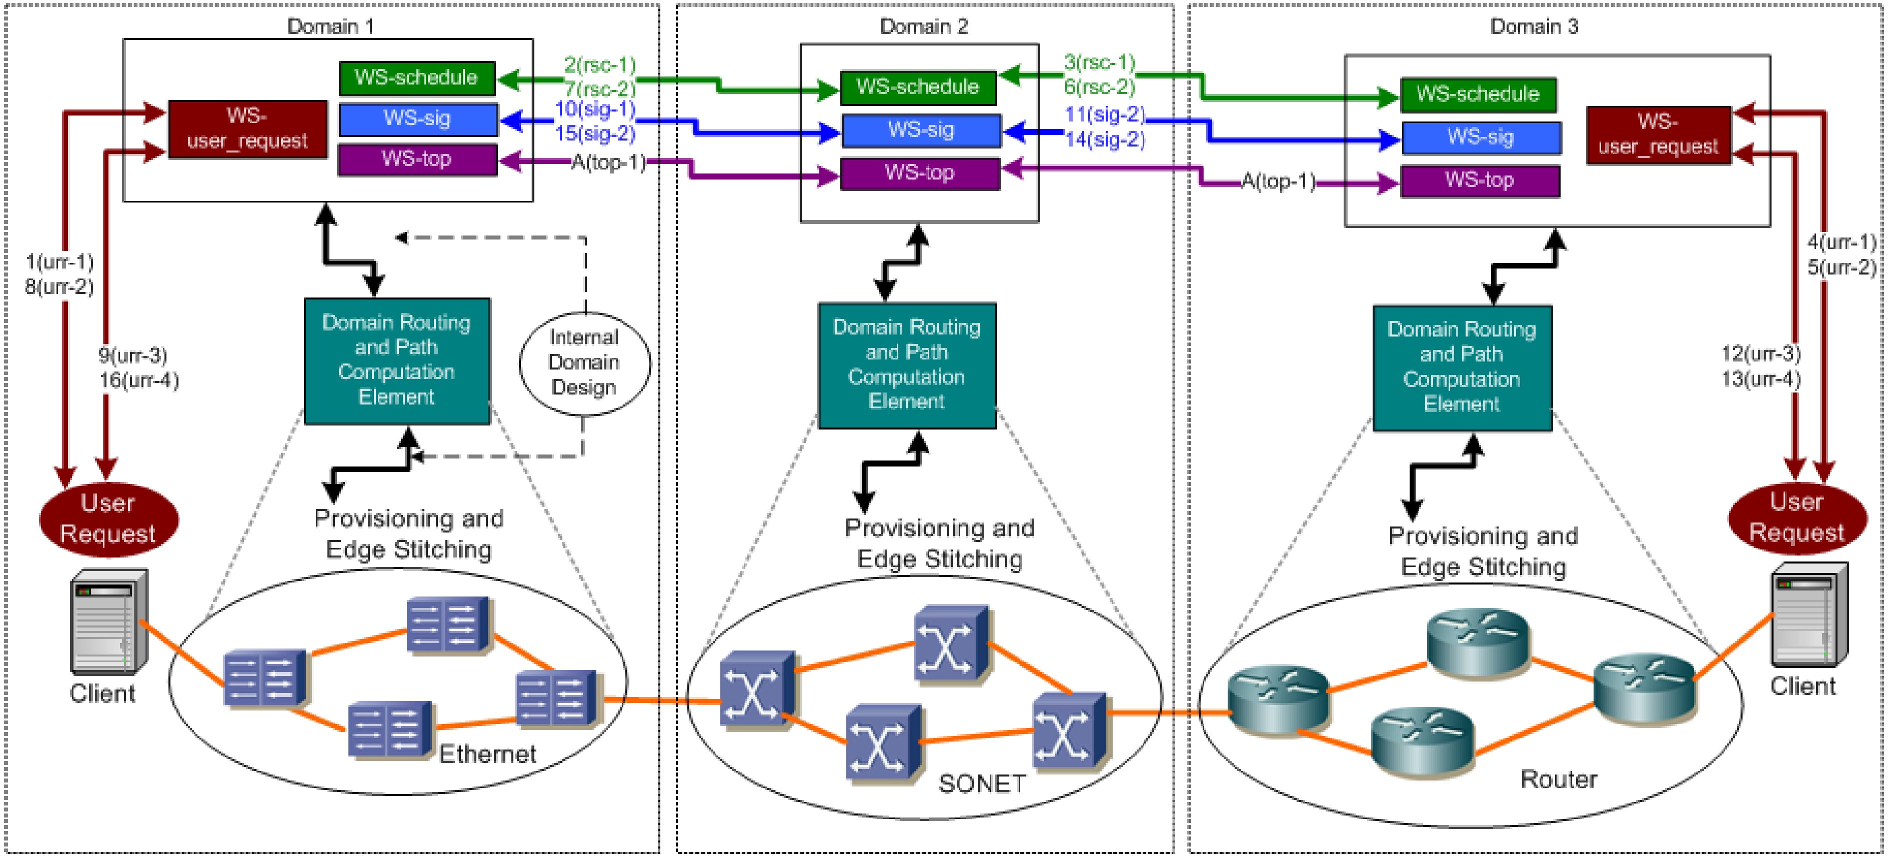

Topology Exchange
top-1: getNetworkTopology

User Request/Response
urr-1: createReservation
urr-2: createReservationResponse
urr-3: createPath
urr-4: createPathResponse

Resource Scheduling
rsc-1: createReservation
rsc-2: createReservationResponse

Signaling
sig-1: createPath
sig-2: createPathReponse

- Inter-domain virtual circuit reservation involves brokers in each domain that present a uniform view of the different domain circuit management mechanisms (e.g. OSCARS in ESnet) (Diagram by Tom Lehman, ISI)

ESnet Virtual Circuit Service Roadmap

- Dedicated virtual circuits
- Dynamic virtual circuit allocation
- Generalized MPLS (GMPLS)

Initial production service

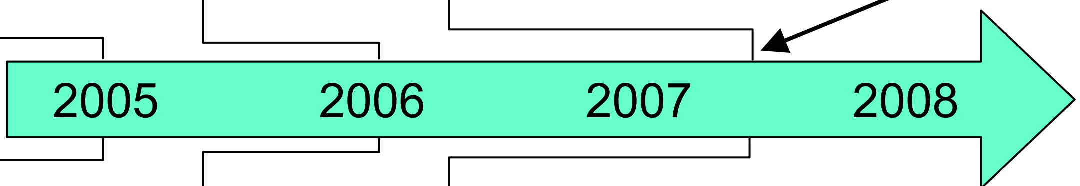

Full production service

- Interoperability between GMPLS circuits, VLANs, and MPLS circuits (layer 1-3)
- Interoperability between VLANs and MPLS circuits (layer 2 & 3)
- Dynamic provisioning of Multi-Protocol Label Switching (MPLS) circuits in IP nets (layer 3) and in VLANs for Ethernets (layer 2)

Monitoring as a Service-Oriented Communications Service

- perfSONAR is a community effort to define network management data exchange protocols, and standardized measurement data gathering and archiving
- Path performance monitoring is an example of a perfSONAR application
  - provide users/applications with the end-to-end, multi-domain traffic and bandwidth availability
  - provide real-time performance such as path utilization and/or packet drop
- Multi-domain path performance monitoring tools are in development based on perfSONAR protocols and infrastructure
  - One example – Traceroute Visualizer [TrViz] – has been deployed in about 10 R&E networks in the US and Europe that have deployed at least some of the required perfSONAR measurement archives to support the tool

➢Federated Trust Services – Support for Large-Scale Collaboration

- Remote, multi-institutional, identity authentication is critical for distributed, collaborative science in order to permit sharing widely distributed computing and data resources, and other Grid services
- Public Key Infrastructure (PKI) is used to formalize the existing web of trust within science collaborations and to extend that trust into cyber space
  - The function, form, and policy of the ESnet trust services are driven entirely by the requirements of the science community and by direct input from the science community
- International scope trust agreements that encompass many organizations are crucial for large-scale collaborations
  - ESnet has lead in negotiating and managing the cross-site, cross-organization, and international trust relationships to provide policies that are tailored for collaborative science
  - ➢ This service, together with the associated ESnet PKI service, is the basis of the routine sharing of HEP Grid-based computing resources between US and Europe

➢ESnet Conferencing Service (ECS)

- A highly successful ESnet Science Service that provides audio, video, and data teleconferencing service to support human collaboration of DOE science
  - Seamless voice, video, and data teleconferencing is important for geographically dispersed scientific collaborators
  - Provides the central scheduling essential for global collaborations
  - ESnet serves more than a thousand DOE researchers and collaborators worldwide
    - H.323 (IP) videoconferences (4000 port hours per month and rising)
    - audio conferencing (2500 port hours per month) (constant)
    - data conferencing (150 port hours per month)
    - Web-based, automated registration and scheduling for all of these services
  - Very cost effective (saves the Labs a lot of money)

➢Summary

- **ESnet is currently satisfying its mission by enabling SC science that is dependant on networking and distributed, large-scale collaboration:**

  "The performance of ESnet over the past year has been excellent, with only minimal unscheduled down time. The reliability of the core infrastructure is excellent. Availability for users is also excellent" - DOE 2005 annual review of LBL

- **ESnet has put considerable effort into gathering requirements from the DOE science community, and has a forward-looking plan and expertise to meet the five-year SC requirements**
  - **A Lehman review of ESnet (Feb, 2006) has strongly endorsed the plan presented here**

References

1. High Performance Network Planning Workshop, August 2002
   – http://www.doecollaboratory.org/meetings/hpnpw
2. Science Case Studies Update, 2006 (contact eli@es.net)
3. DOE Science Networking Roadmap Meeting, June 2003
   – http://www.es.net/hypertext/welcome/pr/Roadmap/index.html
4. DOE Workshop on Ultra High-Speed Transport Protocols and Network Provisioning for Large-Scale Science Applications, April 2003
   – http://www.csm.ornl.gov/ghpn/wk2003
5. Science Case for Large Scale Simulation, June 2003
   – http://www.pnl.gov/scales/
6. Workshop on the Road Map for the Revitalization of High End Computing, June 2003
   – http://www.cra.org/Activities/workshops/nitrd
   – http://www.sc.doe.gov/ascr/20040510_hecrtf.pdf (public report)
7. ASCR Strategic Planning Workshop, July 2003
   – http://www.fp-mcs.anl.gov/ascr-july03spw
8. Planning Workshops-Office of Science Data-Management Strategy, March & May 2004
   – http://www-conf.slac.stanford.edu/dmw2004
9. For more information contact Chin Guok (chin@es.net). Also see
   - http://www.es.net/oscars

[LHC/CMS]
http://cmsdoc.cern.ch/cms/aprom/phedex/prod/Activity::RatePlots?view=global

[ICFA SCIC] "Networking for High Energy Physics." International Committee for Future Accelerators (ICFA), Standing Committee on Inter-Regional Connectivity (SCIC), Professor Harvey Newman, Caltech, Chairperson.
- http://monalisa.caltech.edu:8080/Slides/ICFASCIC2007/

[E2EMON] Geant2 E2E Monitoring System –developed and operated by JRA4/WI3, with implementation done at DFN
http://cnmdev.lrz-muenchen.de/e2e/html/G2_E2E_index.html
http://cnmdev.lrz-muenchen.de/e2e/lhc/G2_E2E_index.html

[TrViz] ESnet PerfSONAR Traceroute Visualizer
https://performance.es.net/cgi-bin/level0/perfsonar-trace.cgi